RC22345 (W0202-056) February 20, 2002
Electrical Engineering

IBM Research Report

The Physical Design of on-Chip Interconnections: Part I: Quantification of Custom Interconnections

Mary Y. L. Wisniewski, Emmanuel Yashchin, Robert L. Franch, Dan Maynard, David Conrady, Giovanni Fiorenza, I. Cevdet Noyan
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Delhi - Haifa - India - T. J. Watson - Tokyo - Zurich

The Physical Design of on-Chip Interconnections: Part I: Quantification of Interconnect Properties

M. Y. L. Wisniewski, E. Yashchin, R. L. Franch, D. Conrady, D. Maynard, G. Fiorenza, I. C. Noyan

M. Y. L. Wisniewski, E. Yashchin, R. L. Franch, G. Fiorenza and I. C. Noyan are with the IBM T.J. Watson Research Center, Route 134, Yorktown Heights, NY. D. Maynard is with the IBM Microelectronics Division in Burlington, VT. D. Conrady is with the IBM Microelectronics Division in Poughkeepsie, NY.

Abstract

Custom interconnect design complements automated route algorithms which do not guarantee the generation of robust, legal routes for all signals in a ULSI design. Intervention with custom interconnections stabilizes a route solution and reduces time required to route follow-on designs with logic and floorplan changes. Implementation of custom interconnects is a highly skilled *art*. In general, the tools available to the designer do not evaluate the *quality* of the inserted custom interconnections, and success is predicated only on the production of a legal design within a pre-specified time. It is desirable to have criteria to select signals for custom interconnections and to evaluate whether or not intervention with custom interconnections for these signals is effective in improving physical properties of these signal routes without adversely affecting physical properties of the remaining signal routes. This paper, the first in a series, presents a self-consistent formalism of intervention with custom interconnections, the addition of custom interconnections to a physical design, and the quantification of overall interconnection quality. The analytical techniques presented in this series of papers can be also incorporated in semi-custom and ASIC designs and may serve as tools to evaluate and improve various route algorithms.

Keywords

Custom interconnect design, custom interconnection, signal density, via density, normalized excess Manhattan length, normalized excess Steiner length, cumulative effectiveness, netlength effectiveness, via effectiveness, stacked via.

I. Introduction

An understanding of the role of interconnections in ultra-large-scale-integrated (ULSI) chip design is important to achieve optimal performance in high-speed microprocessors and has implications for the manufacturability and realization of increasingly complex circuits[1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12]. The use of large numbers of signals in ULSI designs increases design complexity, and the importance of understanding the effects of this increasing complexity has been highlighted by the Semiconductor Industry Association[3]. Moreover, the detailed design of interconnections for these signals also impacts design yield, performance, and power dissipation as well as system cost and information processing ability[4], [13], [14], [15], [16], [17], [18], [19], [20].

Methods to optimize interconnect physical characteristics are desired since these characteristics affect interconnect electrical properties and reliability in addition to chip and system characteristics mentioned above. For this reasons, it is useful to have quantitative criteria to evaluate the value and impact, if any, of *a priori* interconnect design techniques

on interconnect physical properties. Custom interconnect design is one type of *a priori* interconnect design technique. To achieve an optimum design, *a priori* design techniques are incorporated with automated route algorithms in all stages of the design process, including early stages, in contrast with *a posteriori* techniques that perform manipulations of routes in a completely-routed design.

Automated route algorithms that route ULSI designs are complemented by custom interconnect design since automated route algorithms do not guarantee the generation of a stable, legal (that is, violation-free, design-rule-correct, zero unrouted signals) route solution for all design signals. Custom interconnections stabilize a route solution, reduce time to reroute a design, and reduce time to incorporate logic and area changes in current and follow-on designs.

This paper is the first in a series on the subject of physical design of on-chip interconnections. This paper presents a quantitative method for custom interconnect design as well as quantitative criteria to analyze physical characteristics of interconnections in ULSI designs. The analytical techniques presented in this series of papers can be incorporated into semi-custom and ASIC designs as well as custom designs and may serve as tools to evaluate and improve various route algorithms.

In this paper, we present a self-consistent formalism on intervention with custom interconnection design and the quantification of overall interconnect design quality. Included in this formalism are techniques to add custom interconnections to a design in a series of steps (the series of steps can also be referred to as a series of *trials*), to create two-dimensional visualizations of the signal and via density in a design both before and after the addition of custom interconnections, to measure quality of total signal netlength and number of non-redundant vias, and to measure excess length of each signal route compared to a benchmark length for that signal route. Section II reviews characteristics of custom interconnections, Section III describes custom interconnection design, Section IV describes intervention with custom interconnections, and Section V quantifies interconnect physical characteristics. Sections VI and VII describe implications for ULSI design and scaling, respectively.

In each paper in this series, the presented techniques are applied to analyze intercon-

nections in the IBM POWER4 Instruction Fetch Unit. The POWER4 high-performance microprocessor is incorporated in the IBM Enterprise Server pSeries 680.[21], [22], [23] The $415mm^2$ POWER4 chip contains 174 million transistors and is manufactured with IBM's 1.6 Volt, 0.15 μm, dual V_t, Silicon-on-Insulator (SOI) CMOS technology with seven layers of copper interconnections. The power dissipation of the POWER4 chip at 1.6V and 1.3 GHz is 140 *Watts*.[21], [24] Each POWER4 chip contains two microprocessor cores, a shared L2 cache, a directory for the (off-chip) L3 cache, and logic required to connect multiple POWER4 chips. This paper discusses interconnect design in two versions of the POWER4: DD1 (Design-Data-1) and DD2 (Design-Data-2). The term DD1 refers to the initial implementation of a 700MHz (1430ps cycle time) POWER4 chip completed in 3Q 1999.[25] The term DD2 refers to a 1.1GHz (910ps cycle time) POWER4 chip completed in 3Q 2000.[25]

II. What is custom interconnect design?

Custom interconnect design is the process by which signals in a design are wired with *custom interconnections*. In this paper, the term *custom interconnection* refers to a wire that is carefully laid out by hand and is inserted in a physical design. Custom interconnections provide several benefits, including the ability to: (1) obtain a violation-free route solution for signals that cannot be correctly wired by an automated router; (2) decrease *interconnect complexity* by reducing signal density and via density in congested regions; (3) remove timing violations and poor risetimes and falltimes on timing-critical signals; (4) permit easy route iterations after logic and signal changes; (5) iterate timing and routing of all unit-level signals; (6) quick remap of custom interconnection solutions from one design to a similar follow-on design (within $1-2$ weeks, thereby saving considerable designer time on the order of several months).

Custom interconnections are unchanged during the completion of routes for the remaining signals by an automated router. Custom interconnections tend to require careful thought and considerable design time to implement (e.g., six months to implement custom interconnections for approximately 37% of the POWER4 IFU unit-level signals). Implementation of custom interconnections is practical and straightforward once unit constraints on signal routes are specified. Constraints include the unit area, number and location of

unit IO pins, number of unit-level signals, average fanout per unit-level signal, unit occupancy, and fraction of total route length available for unit-level interconnections. The average fanout is a measure of the degree of signal connectivity in the unit. The unit occupancy is the fraction of the unit area reserved for macros that contain devices (transistors).

The term *interconnect complexity* refers to the physical characteristics of interconnections in a design. These physical characteristics include: (1) interconnect length[1] of each signal; (2) the amount of extra length that is required to route a signal relative to a benchmark estimate such as the Steiner length[26] or Manhattan length[27]); (3) the number of non-redundant vias[2] in a signal route; (4) signal density; (5) via density; and (6) number of stacked vias in the signal routes. The term *stacked via* refers to via stacks that are composed of two vias (a *single stacked via*), three vias (a *double stacked via*), four vias (a *triple stacked via*), and so on. To minimize interconnect complexity, custom interconnections are carefully planned for unit-level signals in timing-critical paths and in congested regions to reduce interconnect length and delay and to reduce signal and via density.

Figure 1 shows examples of custom interconnections for signals with fanout=1 and four different types of pins on the connected macros. In Fig. 1(a), a signal is routed on a short straight *m2* wire segment between two *m2* square pins that share the same horizontal track on both macros, as shown in the top and side views. In Fig. 1(b), a signal is routed on a long straight *m4* wire segment between two *m4* square pins that share the same horizontal track on both macros. In Fig. 1(c), a signal is routed on a straight *m4* wire segment between an *m4* pin on one macro and through a *v3* via to an *m3* horizontal bar pin on the other macro. For macros with pins on different metal layers and different relative horizontal tracks, Fig. 1(d), shows a signal routed on a long straight *m4* wire segment from an *m4* pin to a *v3* via and *m3* wire segment and then to a *v2* via and *m2* wire segment and *m2* pin.

Different types of custom interconnections route signals between custom macros and

[1]In this paper, the interconnect length of each signal is obtained from the physical design by calculating the sum of the lengths of metal segments that route a signal to all connecting pins.

[2]In this paper, the number of vias in a signal route is obtained from the physical design by calculating the sum of vias that connect all the metal segments that route a signal to all connecting pins.

random logic macros (*rlms*). Figure 2 shows examples of custom interconnections that route signals between three custom macro abstracts and an rlm abstract: (a1, a2, a3) completely-routed minimum-width routes (these routes, for example, route signals over the predecode macros and in the address stack); (b) completely routed routes with double-wide routes and minimum-width routes (these wires route signals in the wiring bay above the instruction cache stack); (c) partially routed minimum-width routes (these wires route signals with large fanout in the address stack); and (d) partially routed wide routes (these wires route buses over the instruction cache), as shown in Fig. 2(a), Fig. 2(b), Fig. 2(c), and Fig. 2(d), respectively. Partial custom interconnections are useful in cases in which the complete specification of a route is not required. For example, partial routes are advantageous to route long sections of signals with large fanout or to route long straight point-to-point signals; an automated router then completes the unrouted portion of the signal. To help meet the cycle-time requirement, timing-critical signals are routed with wide metal to reduce interconnect delay.

Custom interconnections also route buses between two or more custom macros. Figure 3 shows fully-routed minimum-width custom interconnections (such as *a1-a3* in Fig. 2) that route two 16-bit buses between dataflow macros. In this example, the dataflow stack is composed of three 16-bit-wide dataflow macros. The first set of 16 minimum-width wires routes a 16-bit bus on minimum-width *m2* wires between aligned sets of *m2* pins on two adjacent macros (namely, the middle macro and lowest macro). The second set of 16 minimum-width wires routes a second 16-bit bus on minimum-width *m4* wires that connect aligned sets of *m3* bar pins on the upper-most macro and the lowest macro. This bus crosses the middle macro in a track that contains no *m4* in the middle macro, which permits the straight connection on *m4* between the lowest macro with the upper-most macro. As shown in the figure, the upper-most macro has a one-bit-wide extension on the right-hand side and another one-bit-wide extension on the left-hand side for control logic. The middle dataflow macro has a two-bit-wide extension on the right-hand-side for control logic; the bottom dataflow macro does not have extensions. The two two-bit-wide columns and horizontal *m3* pins represent clock bays and clock control pins, respectively. For the case of real IFU dataflow abstracts, each dataflow macro is 64-bits wide and each

4-bit-wide clock bay is located between bit 15 and 16 and 47 and 48, respectively.

A. Application to the POWER4 IFU

The POWER4 IFU contains approximately 5.9 million transistors (including 4 million array transistors). The total length of copper interconnect in the IFU exceeds 16 meters for all unit-level signals including power and ground. Excluding power and ground, the total interconnect length in the IFU exceeds 5 meters. There are slightly more signals (9283) in the DD2 IFU compared to the number (9253) in the DD1 IFU. In the case of the IFU, as discussed in a later section, intervention with custom interconnections permits the IFU physical design area to be reduced approximately $0.2mm^2$ (1.7%) from DD1 IFU to DD2 IFU yet remain routable despite the overall decrease in available length for interconnections. The area decrease can decrease total chip area and increase the number of manufacturable chips per wafer. Minimum-length custom interconnections with fewest possible non-redundant vias also tend to increase chip yield[4]. This result is important since chip yield is a serious concern for large chips.

III. Custom interconnect design

The previous section described custom interconnect design and presented a few examples of custom interconnections. In this section, we describe the process of custom interconnect design: (1) to remove route violations generated with an automatic router, (2) to route complex designs, and (3) to reduce interconnect complexity. This method addresses the current problem that automated route algorithms does not guarantee the generation of a robust, legal route solution (that is, violation-free, design-rule-correct, zero unrouted signals) for all signals in a ULSI design. Currently, implementation of custom interconnections is a highly skilled *art*, and there is a need for tools to help designers complement automated route algorithms with intervention with custom interconnections.

A. To remove route violations generated with an automated router

For designs that cannot be routed by an automated router without generating violations, custom interconnections replace signal routes to remove these violations.

A.1 Application to the POWER4 IFU

Figure 4 shows that the automated router is unable to route the design without also adding design-rule violations and electrical shorts in the signal routes ($N_c = 0$). The figure also shows that the optimum point for the routes in the DD1 design occurs when the design contains custom interconnections for 37% of all unit-level signals. This point occurs when all design rule violations and electrical violations are eliminated.

B. To route complex designs

Complex designs exhibit high values of design occupancy O. The occupancy O is defined to be the fraction of the unit area A that is occupied by logic macros that contain devices.

Second, complex designs contain regions with a relatively small amount of available route length compared to the total estimated signal length. For example, a design is complex when the number of signals is on the order of the number of available route tracks. For example, a complex design results when the number of (vertical) signals in per bit approaches the number of available (vertical) metal tracks (per bit) over a macro and when the number of (horizontal) control signals (per bit) approaches the number of available (horizontal) wiring tracks. The amount of route length that is available to route unit-level signals with minimum-width wire on each of the available route layers is also measured.

Complex designs also contain many timing-critical signals, such as address buses or instruction buses that route dataflow logic signals throughout the design. To fix the delay along the paths that contain these signals, custom interconnections route these buses in order to fix the route solutions, and therefore delay, for these signals for the duration of the project. Remaining non-timing critical signals can be routed during subsequent iterations of routing and timing with an automated router.

B.1 Application to the POWER4 IFU

Figure 5(b) shows the IFU occupancy O as a function of time. For the first 16 months, values of the macro areas are estimated by the logic designers. After 17 months, circuit designers start to complete the physical design of the macros to establish each macro area. As these results replace the earlier estimates, the unit occupancy O approaches

73%. Figure 5(b) shows that the unit occupancy increased monotonically from 17 to 23 months into the project. This trend occurs since approximately 1000 buffers were added to the design to: (1) fix signals with risetimes or falltimes larger than 330ps; (2) buffer timing-critical unit-level signals with lengths greater than 1.4mm; and (3) buffer signals that connect macros to unit-level IOs. In the latter case, buffers and inverters were placed within 0.7mm of unit IO pins. Figure 5 shows the IFU unit area A as a function of time. Early estimates by IFU logic designers placed the unit area size at approximately $16mm^2$, except at 6 months when the estimate for A was temporarily tightly squeezed. By the time of the DD1 snapshot at 22 months, A decreased 12.5% to $14mm^2$ and again slightly 1.9% for the DD2 snapshot at 33 months. The goal of this area decrease was to open additional chip area for other functional units. Figure 5 shows that the values of A (a) and O (b) change only slightly over time, which indicates that the preliminary macro and unit area estimates have been modeled accurately.

For the case of the IFU, 8 to 10 tracks are available over each bit of a typical macro to route up to 8 to 10 signals with minimum-width wire (since 8 to 10 tracks of *m4* are free over a typical dataflow macro). In addition, 4 horizontal tracks are typically available per bit to route 4 signals per bit across each dataflow macro (since approximately 70% = 7/10 of free *m3* tracks and 10% = 1/10 of *m5* tracks are available over a typical dataflow macro). Additional horizontal tracks are available in the region between macros to route signals horizontally; however, this inter-macro spacing is reduced as much as possible to conserve chip area. Typically, the inter-macro spacing is less than $18\mu m$, which permits approximately 10 to 20 signals to be routed horizontally between macros. Figure 6(a) shows the maximum route length and the available route length for unit-level signals as a function of time. Prior to 15 months, the values for the route lengths were determined from macro and unit area estimates provided by logic designers and from estimates of macro metal usage. Prior to 16 months, 100% of *m1* and *m2*, 70% of *m3* and 0% of *m4* were assigned to macros. At 16 months, circuit designers begin completing the macros' physical design with areas that tend to be smaller than the estimated values, as shown in the figure. By month 23, the DD1 physical design is complete. After an additional 1.9% area shrink is imposed from DD1 to DD2, the maximum available route length and

total available route length both decrease slightly. Lower-level metal (upper-level metal) represents approximately $20\% - 30\%$ ($70\% - 80\%$) of all available interconnect, as shown in Fig. 6(b).

Figure 7(a) shows the available route length on each of five metal layers as a function of time. The total available route length is composed of contributions from lower-level-metal layers (*m1-m2*) and upper-level-metal layers (*m3-m5*). The fraction of L_{max} that is available to route minimum-width unit-level signals is shown in Fig. 7(b). This figure shows that the length of lower-level metal that is available to route unit-level signals is less than the length of upper-level metal, which occurs because most tracks on lower-level-metal layers are reserved by macros for internal macro routes. Standard design practice recommends that long unit-level signals should be routed on upper-level metal since these layers have lower resistivity. Short signals that connect aligned macro pins can be routed on straight segments of lower-level metal (e.g., *m2* wire to connect *m2* pins). The total fraction of L_{max} available for routing minimum-width unit-level signals remained nearly constant throughout the project, with $35\% - 40\%$ available at 15 months and 40% and 37% available in DD1 and DD2, respectively. From these figures, we see that the amount of available lower-level metal and upper-level metal are quite close to the values estimated from a preliminary floorplan and macro blockage estimates, which shows that the early values have also been accurately modeled.

C. To reduce interconnect complexity

One measure of interconnect complexity is the total interconnect length L_T in all signal routes at a specified level of the hierarchy. Another measure of interconnect complexity is the signal density and via density in the design. Congested regions tend to exhibit a relatively high value of signal density and via density compared to that in neighboring regions.

C.1 Application to the POWER4 IFU

For the IFU, $L_T = 5.5m$ in DD1 ($L_T = 5.4m$ in DD2). This value can be compared to: (1) the width of a space allocated to a wire; (2) the unit area that contains the unit-level routes; or (3) the unit height in order to obtain three relative measures of the

complexity of the interconnections, as applied by Keyes to a chip manufactured in 1981 [4]. Compared to the 1981 chip[4], which we take here to be a reference example, the total length of interconnections in the IFU (excluding power and ground interconnections) is approximately 1.7×10^7 times the width of a space allocated to a wire, compared to approximately 6×10^5 for all routes on the reference chip[4]; $A_{int} = 0.9A$, compared to approximately 1 times the reference chip area for all the reference chip routes[4]; and L_T is approximately 1100 times the unit height, compared to 800 times the reference chip edge[4].

The IFU contains three congested regions which are discussed in detail in Section V. The first region is located over the instruction cache stack, as shown on the left-hand side of Figs. 8. Forty-two buses, including eight 32-bit instruction buses, are routed vertically over the instruction cache stack. Buses associated with these instructions are routed horizontally in the wiring bay above the instruction cache stack on the left-hand-side. The second region is the address stack, shown on the right-hand side. Forty-eight buses, including many 64-bit address buses are routed between macros in the address stack. The third region is located over the nine *rlms* that implement the predecode function in the IFU. These *rlms* are located along the bottom of the left-hand-side of the unit at the bottom of the instruction cache stack. Forty-two buses are routed horizontally to connect the central macro with four macros to the left and four macros to the right.

D. To route timing-critical signals

Custom interconnections route signals that pass through non-congested regions and replace routes for timing-critical signals that have been routed by the automated router with excessively long routes. Intervention with custom interconnections also corrects signal routes with poor risetimes and falltimes. Custom interconnections are also desirable to impose the same route solution for all signals in a bus or group of buses in order to equalize interconnect delay for these signals.

IV. Method to insert custom interconnections

In the previous section, we described designs for which intervention with custom interconnection is useful. In this section, we describe a method to identify signals with excess

route length, identify congested regions, and insert custom interconnections.

A. Identify signals with excess route length

The first step is to identify signals to target for custom interconnections. In this step, we to route the design that contains custom macros and *rlm* macros either with an automatic router or with a mixture of automatically routed signals and custom interconnections. This design will be analyzed even if some routes contain violations such as electrical shorts or geometry violations.

In the routed design, the netlength L, Steiner length L_S, and Manhattan length L_M are measured for each signal route. To describe the route quality of each individual unit-level signal, we introduce two normalized quantities: (1) the normalized excess Steiner length *NESL* given by the expression,

$$NESL = \frac{\Delta L_S}{L_S} = \frac{(L - L_S)}{L_S}, \tag{1}$$

and (2) the normalized excess Manhattan length *NEML* given by the expression,

$$NEML = \frac{\Delta L_M}{L_M} = \frac{(L - L_M)}{L_M}. \tag{2}$$

For each signal, L_S represents the signal Steiner length that is measured with a Steiner algorithm such as that in IBM's floorplanning tool Hierarchical Design Planner (HDP)[26]. The Manhattan length L_M is the sum of the maximum horizontal and maximum vertical separations of all pins that connect each signal[27]. All unit-level signal routes are then ordered in two lists of decreasing values of *NESL* and *NEML*. The first (second) list shows each signal route and its normalized excess Steiner (Manhattan) length in order of decreasing value of *NESL* (*NEML*). We observe that bus signal routes with large values of *NESL* and *NEML* exhibit similar (large) values of *NESL* and *NEML* as other signals that are contained on the same bus. As a result, the lists can be reduced into two ordered lists of buses with decreasing value of *NESL* and *NEML*, respectively. Next, each bus with large *NESL* and *NEML* in the ordered lists of *NESL* and *NEML* is then located in the design; buses that occupy approximately the same physical location are considered to occupy the same physical region.

B. Identify congested regions

The next step is to identify congested regions in a design. In this step, we generate two two-dimensional contour maps of the signal density n_s and the via density n_v of the routed design to visualize the signal and via distributions. A visual examination of the density maps indicates regions with high n_s and high n_v compared with densities in neighboring regions. Regions with both high n_s and n_v compared to the values in surrounding regions are referred to as *congested regions* and are visualized with two-dimensional color contour maps of n_s and n_v.

B.1 Application to the POWER4 IFU

Figure 9 shows a series of two-dimensional contour maps of n_s in each $35.3\mu m \times 35.3\mu m$ area in an unrouted floorplan. Figure 9(a) shows uniform $n_s = 2$ in the unrouted design due to the power and ground grid, except at the clock pins, where $n_s = 3$. Figure 9(b) and Fig. 9(c) show n_s and n_v, respectively, in a completely routed floorplan. The IFU contains three congested regions, and an automated router is not able to route these regions without generating violations. These regions are marked 1, 2, and 3 in Fig. 8. The number in parenthesis near each region indicates the number of buses identified with high values of *NESL* and *NEML*. Region 1 is the wiring bay above the instruction cache stack, Region 2 is located over the dataflow stack, and Region 3 is located below the instruction cache stack. Four non-congested regions (Regions $4-7$) contain signals targeted for custom interconnections. Region 4 is located over the buffers that repower core-level signals, Region 5 is located over the instruction cache, Region 6 is located at the botton of the dataflow stack near unit IOs, and Region 7 is located above Region 3. Regions 2, 4, and 6 are located on the right-hand side of the unit; Regions 2 and 6 are located in the stack of dataflow macros. Region 2 is congested because the number of bus signals in this region (i.e., 10 bus signals per bit) is equal to the number of available tracks (10); this region also has high values of n_s and n_v as shown in Figs. 9(b) and (c). The automated router is unable to find one of the few possible solutions, and as a result, this congested region contains route violations. Region 6 is uncongested and contains short signal routes that connect macro pins with IO pins. Regions 1, 3, 5, and 7 contain buses that are located on

the left-hand side in the instruction cache stack. Region 1 is a wiring bay that is occupied by address buses that connect macros near the top of the instruction cache with IO pins at the upper left boundary and to dataflow macros located on the right-hand side. For this reason, the signal routes jog horizontally and vertically and are designed with minimum-width routes to occupy minimum vertical space. A few signals in this bay are routed with wide metal to eliminate poor slews. As a result, this region is also highly congested, as shown in Fig. 9(b). Regions 3 and 7 contain horizontally routed buses. Region 3 is uncongested and consists of sparsely spaced horizontal buses routed on *m3*. Region 7 consists of closely-spaced horizontal buses (i.e., up to 5 bus pins per bit) and are routed over macros that occupy 100% of the *m3* horizontal wiring channels. In this case, to make this region routable, the unit has obtained 4 additional *m5* tracks per bit (for a total of 5 *m5* tracks per bit since 1 *m5* track per bit is available over the whole floorplan) to route bus signals with minimum-width horizontal *m5*. With these constraints, the automated router is still unable to route the design without violations. Figures 9(b) and (c) show that n_s and n_v in these regions is larger than in surrounding regions. In addition, *NESL* and *NEML* for each signal in Region 3 is among the largest (on the order of 0.9). Therefore, buses in Region 5 are targeted for custom interconnections.

C. Identify signals to target for custom interconnections

The next step is to compare the locations of congested regions with locations of regions that contain signals with large values of *NESL* and *NEML*. Regions that are both congested and contain buses with large values of *NESL* and *NEML* are regions in which buses are targeted for custom interconnections.

D. Insert custom interconnections

Custom interconnections are inserted in the design for those bus signals that exhibit high values of *NESL* and *NEML*, as discussed above. Custom interconnections are added for buses with smaller values of *NESL* and *NEML* depending on whether the bus pins connect macros in a dataflow stack, whether the bus pins connect macros in a horizontal *rlm* stack, or whether the bus pins connect dataflow macros in different stacks. After custom interconnections are added for each bus, the remaining signals are routed and

checked for electrical shorts and geometry violations. This procedure is iterated until no unit-level signals have large values of *NESL* and *NEML* above a limit specified by the designer, and no route violations are present.

Custom interconnections for signals in a bus are created according to the same procedure; these signals are routed on the same metal layer in the same relative horizontal or vertical route track with the same netlength and number of vias. Individual control signals routes are inspected before a complete custom interconnection is added. For example, custom interconnections can route control signals in the case in which signal pin locations are not expected to change location in a physical design cycle. For control signals involved in many logic changes and control signals with slew violations, the automatic router is instructed to route critical control signals early in the automatic route process.

For each trial i, the total number of signals with custom interconnections N_c^i and the number of additional signals with custom interconnections ΔN_c^i are obtained, where ΔN_c^i is given by the expression,

$$\Delta N_c^i = N_c^i - N_c^{(i-1)}. \tag{3}$$

Information provided by the dimensions and metal blockages of the *unit abstract* and *macro abstracts* determines the amount of wire available to route unit-level interconnections. Constraints include the unit area A, number of IO pins, number of signals N, average signal fan-out, unit occupancy O, total length L_{avail} available to route unit-level signals, and fraction of the maximum route length available to route unit-level signals.

We set the quantity L_{avail} equal to the total length available to route unit-level signals on the five available metal layers. L_{avail} is calculated assuming minimum-width minimum-space routes for all signals. This assumption holds for the case of the DD1 IFU and DD2 IFU, for which 97% of the unit-level signals are minimum-width and minimum-space (the remaining 3% signals are routed with width equal to double or triple the minimum width). L_{avail} is calculated according to the following procedure: First, the maximum length of all minimum-width route channels (including power and ground) on all the available route layers is calculated. Second, the total length of blocked channels (i.e., obstructions) is measured on all metal layers in all macros. Next, the total length of all minimum-width channels blocked by power and ground routes is measured. Then, the total length of

all blocked minimum-width route channels is obtained by summing the total length of channels blocked by macros and by power and ground. The blocked channels are not available to route unit-level signals. L_{avail} is obtained by subtracting the total length of all blocked channels from the maximum length of all route channels.

D.1 Application to the POWER4 IFU

Table I shows an overview of the signals that are routed with custom interconnections in the DD1 IFU. A total of 133 buses (3390 bus signals) and 11 control signals are routed with custom interconnections (either complete custom interconnections or partial custom interconnections). The table shows the number of bus signals N_{bus}, number of control signals N_{cntl}, and number of buses inserted in each region $N_{reg1}, ..., N_{reg7}$. The total number of buses inserted in all regions is N_{regtot}, where $N_{regtot} = \sum_{i=1}^{7} N_{reg_i}$. For the IFU, the congested regions have not been routed according to the highest degree of congestion (in this scenario, Region 3 would be routed first, followed by Region 2, then Region 1, followed by the last four regions). Instead, due to project constraints and degree of macro specification, the regions have been routed in a slightly different order (namely, half of Region 2 is routed first, then Region 1, then Region 3). Region 2 is routed first because logic in the dataflow macros and therefore their abstracts and bus pin locations are stabilized earliest in the project schedule. As shown in Table I, slightly fewer than half of the buses in Region 2 are routed initially in order to verify that a route solution does exist with available *m4* route resources.

Table II shows the physical design constraints on the IFU unit-level interconnections. A detailed distribution of the unit-level signal fanout is shown in the following table Table III. Table II shows that the decrease in the unit area reduces the total available route length 1.9% from $31m$ in DD1 to $29m$ in DD2. The available route fraction is obtained by taking the ratio of L_{avail} to the maximum total length of the minimum-width route channels. Table II shows that the portion of all the minimum-width tracks available to route unit-level signals is reduced from 34% in DD1 to 31% in DD2.

Of the seven copper metal layers available to route the POWER4 chip, the five lowest metal layers are available to route IFU signals. These layers are: metal 1 (*m1*), the lowest route layer; *m2*, *m3*, *m4*, and *m5*. Each metal layer is routed in the plane parallel to the

silicon layer and perpendicular to the metal layer above and below. In this paper, the phrase *upper-level metal* refers to layers *m3*, *m4*, and *m5*; the phrase *lower-level metal* refers to layers *m1* and *m2*.

V. Method to quantify interconnect quality

In the previous section, we described a method to identify signals to route with custom interconnections and to insert custom interconnections for those signals. In this section, we analyze the physical elements of an interconnection and present techniques to quantify interconnect quality. We present techniques to quantify several interconnect characteristics: (1) metal usage, (2) signal density, (3) via density, and (4) complexity.

A. Quantify interconnect metal usage

One measure of the quality of interconnections is the amount of upper-level metal that routes long and short signals. Standard design practice recommends that long signals are routed on upper-level metal to reduce interconnect delay, and that short signals are routed on lower-level metal.

A.1 Application to the POWER4 IFU

Figure 10 shows (a) the average fraction f_s of upper-level metal in short signal routes with $L \leq 0.7mm$ and (b) the average fraction f_l of upper-level metal in long signal routes with $L \geq 0.7mm$ as a function of N_c in the DD1 IFU. Figure 10(a) shows that f_s decreases to approximately 92% when all targeted signals have been routed with custom interconnections. The fraction f_s in non-custom interconnections remains largely unaffected by the intervention with custom interconnections. Figure 10 shows that, as N_c is increased, f_l increases for custom interconnections and decreases slightly for non-custom interconnections.

B. Quantify interconnect signal density and via density

An overview of the physical properties of interconnections in a design is obtained from two-dimensional visualizations of n_s and n_v as a function of N_c in the design. Next, a series of two-dimensional visualizations are obtained for the design routed with different

values of N_c. From each contour map, horizontal slices (at constant height) and vertical slices (at constant width) for each N_c are obtained.

B.1 Application to the POWER4 IFU

Figures 11(a)- 11(c) show a series of two-dimensional visualizations of n_s for three different values of N_c: 0 (0% of all unit-level signals), 1853 (20.0%), and 3401 (36.8%)). The density maps in Fig. 11 show n_s for custom interconnections in each $35.3\mu m \times 35.3\mu m$ area. The color blue indicates regions with low n_s. Unit-level clock pins are visible in Fig. 11(a) since the clock signal and power and ground grid are routed before all other signals[28]. Clock routes consist of straight *m4* segments that connect *m4* macro clock pins to IO *m5* clock pins. Eight instruction buses are routed over the instruction cache on the left-hand side of Figs. 11(b) and 11(c). Vertical address buses are routed over the address stack on the right-hand side of Figs. 11(b) and 11(c). Figure 11(a') shows n_s for the design shown in Fig. 11(a) routed with an automated router (WarpRoute).

All signals shown in Fig. 11(a') are routed with the automated router (WarpRoute) except for the power, ground, and clock signals. The color red indicates the region with highest n_s; this region contains horizontal routes for eight instruction buses. In this region, n_s exceeds 140 signals per $35.3\mu m \times 35.3\mu m$ area. Figures 11(b) and 11(c) are also routed with the automated router. Relative changes in n_s between two routed designs that contain custom interconnections and that in Fig. 11(a) which contains no custom interconnections are shown in Figs. 11(a') and 11(b'), respectively. The color green indicates regions with reduced n_s compared with regions in Fig. 11(a'). The color blue indicates regions in the instruction cache stack and address stack in which n_s is substantially reduced. These reductions occur because N_c is large in these regions. The red and yellow regions have increased values of n_s. Figures 11(b') and 11(c') show that n_s is reduced in regions that contain custom interconnections.

Figure 12(a) shows a two-dimensional visualization of n_v for the same interconnections shown in Fig. 11(a'). n_v is the number of vias in signal interconnections (excluding power, ground and the clock signal) per $35.3\mu m \times 35.3\mu m$ area. The two red regions have $n_v \geq 100$ per $35.3\mu m \times 35.3\mu m$ area. These two regions are located in the instruction cache stack on the left-hand side of Fig. 12(a) and in the address stack on the right-hand side of Fig. 12(a).

Figures 12(b) and (c) show visualizations of the differences in via density between n_v in Fig. 12(a) and n_v measured for interconnections in Figs. 11(b') and (c'). The color green indicates regions with reduced n_v compared with Fig. 12(a). The color blue indicates regions in which large numbers of custom interconnections are added in the instruction cache stack on the left-hand side of Fig. 11(b) and Fig. 11(c) and the address stack on the right-hand side of Fig. 11(b) and Fig. 11(c). The blue regions exhibit substantially reduced n_v. In Figs. 12(b) and (c), the colors red and yellow indicate regions with larger n_v. These figures show that n_v is reduced in regions that contain custom interconnections. The decrease in the via number in the route process reduces the number of channels blocked by vias. As a result, more channels are available to route other signals, such as signals previously routed with the automated router.

Figure 13 shows visualizations of (a) n_s and (b) n_v per $35.3\mu m \times 35.3\mu m$ area (i.e., an area that is 4-bits wide $\times$ 4-bits wide) in a routed DD1 IFU with $N_c/N = 0\%$ custom interconnections. The dashed lines indicate locations of the cross-sections of n_s and n_v shown in Figs. 14 and 15. These two figures show three cross-sections of (a) n_s and (b) n_v per $35.3\mu m \times 35.3\mu m$ area for three different amounts of custom interconnections inserted in the design prior to the completion of the remaining signal routes by the automated router (0% from Fig. 13, 20%, 38.6%). Figure 14 shows (a) n_s and (b) n_v vias per $35.3\mu m \times 35.3\mu m$ area across the IFU width for three values of the height Y (a1,b1 at $Y = 3669.12\mu m$; a2,b2 at $Y = 3104.64\mu m$; a3,b3 at $Y = 352.8\mu m$). Figure 15 shows (a) n_s and (b) n_v per $35.3\mu m \times 35.3\mu m$ area across the IFU height for two values of the width X (a4,b4 at $X = 1658.16\mu m$; a5,b5 at $X = 2963.52\mu m$; a6,b6 at $X = 3326.4\mu m$). Peaks in the cross-sections correspond to locations of congested regions. The figures show that custom interconnections decrease values of n_s and n_v in congested regions compared with values of n_s and n_v in a design with $N_c = 0$.

C. Quantify quality of design routes

A measure of the overall quality of interconnections in a design is the total netlength L_T and total number of vias v_T required to route all N signals in the design, where L_T

and v_T are given by the expressions,

$$L_T = \sum_{i=1}^{N} L(i), \tag{4}$$

and

$$v_T = \sum_{i=1}^{N} v(i), \tag{5}$$

where $L(i)$ and $v(i)$ are the length and number of vias, respectively, in each i^{th} signal. L_T and v_T can be measured as a function of N_c in the design to determine whether L_T and v_T are reduced, as desired for improved interconnect quality.

Another measure of design quality is the number of *stacked vias* in a design, where the term *stacked via* refers to via stacks that are composed of two vias (a *single stacked via SSV*), three vias (a *double stacked via DSV*), four vias (a *triple stacked via TSV*), and so on.

C.1 Application to POWER4 IFU

Figure 16 shows L_T as a function of N_c (lower abscissa) and $\frac{N_c}{N}$ (upper abscissa) in the DD1 IFU. Figure 16(a) shows that the total length of custom interconnections increases with N_c, and Fig. 16(b) shows that the total length of targeted custom interconnections decreases with N_c. For $N_c = 0$, Fig. 16(a) and Fig. 16(b) show that $L_T = 0$ for custom interconnections, since the physical design contains no custom interconnections, and that $L_T = 2.2m$ for signals targeted for custom interconnections. For $N_c = 3401$ ($N_c/N = 0.37$), L_T for custom interconnections equals $2.1m$, which is a reduction of 5% compared with the $2.2m$ required for $N_c = 0$. Figure 16(c) shows the total route length L_T of the remaining $N_r = 5852$ signals as a function of N_c and N_c/N. The total route length for these signals decreases from a maximum of $3.38m$ for $N_c = 0$ to $3.37m$ for $N_c = 3428$. The total route length L_T for all $N = N_c + N_t + N_r$ unit-level signals is shown in Fig. 16(d). This figure shows that L_T for N unit-level signals decreases from a maximum of $5.51m$ for $N_c = 0$ to a minimum of $5.49m$ for $N_c = 3401(N_c/N = 0.37)$. For v_T in the IFU, Figure 17 shows v_T as a function of N_c ($\frac{N_c}{N}$, upper abscissa). Fig. 17(a) shows that v_T for custom interconnections increases with N_c, and Fig. 17(b) shows that v_T for targeted custom interconnections decreases with N_c. For $N_c = 0$, Fig. 17(a) and Fig. 17(b) show

that $v_T = 0$ for custom interconnections, since the physical design contains no custom interconnections, and that approximately 23000 vias are in signals targeted for custom interconnections. For $N_c = 3401$ ($N_c/N = 0.37$), these figures show that v_T for the custom interconnections is approximately 11000, which is reduced 52% compared with v_T for $N_c = 0$. These figures show that $v_T = 0$ for targeted custom interconnections when $N_c = 3401$. Figure 17(c) shows the total number of vias v_T of the remaining $N_r = 5852$ signals as a function of N_c and $\frac{N_c}{N}$. This figure shows that the vias in the remaining signal routes are reduced 2% from approximately 43500 for $N_c = 0$ to approximately 42600 for $N_c = 3401$. The total number of vias v_T in all $N = N_c + N_t + N_r$ unit-level signals is shown in Fig. 17(d). This figure shows that v_T for the N signals decreases 20% from a maximum of approximately 66500 for $N_c = 0$ to a minimum of 53500 for $N_c = 3401$ custom interconnections ($\frac{N_c}{N} = 0.37$).

Table IV shows *SSV*, *DSV*, and *TSV* for the DD1 IFU for each trial i as N_c is increased in the design. The table shows that custom interconnections reduce *SSV* by 28% from 6814 to 4923, *DSV* by 25% from 878 to 657, and *TSV* slightly from 18 to 12.

D. Quantify complexity of each interconnection

The complexity of each individual interconnection can be quantified by measuring several physical properties: (1) the length L of each signal route; (2) the normalized excess Steiner length *NESL* given by Eqn. 1; (3) the normalized excess Manhattan length *NEML* given by Eqn. 2; (4) and the via number v per signal route. The total interconnect length L_T and total number of vias v_T for all N unit-level signals are obtained by adding the contributions for each signal route as specified by Eqns. 4 and 5. For each signal, $NESL$[3] measures how well the actual wire length approximates the length estimate provided by the Steiner algorithm[26], as described by Eqn. 1. As the wire length L approaches L_S, the excess Steiner length approaches zero. For signals with large L, which can occur if the automated router experiences difficulty finding a route solution, ΔL_S can become large. Similarly, for each signal, the excess Manhattan length ΔL_M measures how closely the wire length approximates the Manhattan length[27]; ΔL_M approaches zero as L decreases to L_M, and ΔL_M becomes large when L exceeds L_M.

[3]We are working on an equivalent description for vias.

We represent the total excess Steiner length and total excess Manhattan length with the quantities $TESL$ and $TEML$, respectively. These quantities are given by the expressions,

$$TESL = \Delta L_{TS} = \sum_{i=1}^{N} \Delta L_S(i), \tag{6}$$

and

$$TEML = \Delta L_{TM} = \sum_{i=1}^{N} \Delta L_M(i), \tag{7}$$

respectively, where $\Delta L_S(i) = L(i) - L_S(i)$ and $\Delta L_M(i) = L(i) - L_M(i)$, and both sums are taken over all N signals in the design. $L(i)$, $L_S(i)$, and $L_M(i)$ represent the length, Steiner length of the i^{th} signal, and Manhattan length of the i^{th} signal, respectively.

Two additional measures of complexity of routed interconnections are the average normalized excess Steiner length and the average normalized excess Manhattan length, where each average is taken over a specified group of signals. For example, this group of signals could consist of the entire set of N signals in the design or a set of signals $N_{group} \leq N$ that passes through any of the identified congested regions in the design. We represent the average normalized excess Steiner length and average normalized excess Manhattan length with the quantities $\langle NESL \rangle$ and $\langle NEML \rangle$, respectively, as given by the expressions,

$$\langle NESL \rangle = \frac{\sum_{i=1}^{N_{group}} NESL(i)}{N_{group}}, \tag{8}$$

and

$$\langle NEML \rangle = \frac{\sum_{i=1}^{N_{group}} NEML(i)}{N_{group}}, \tag{9}$$

respectively, where each sum is taken over all of N_{group} signals in the specified group and where $NESL(i)$ and $NEML(i)$ are the $NESL$ and $NEML$, respectively, of the i^{th} signal. In the case that $\langle NESL \rangle$ or $\langle NEML \rangle$ is to be calculated for a group of signals that passes through a specified region, then N_{group} is set equal to the number of signals in that region; in the case that $\langle NESL \rangle$ or $\langle NEML \rangle$ is to be calculated for all signals, then we set $N_{group} = N$.

A corresponding measure of via complexity can be expressed as the average number of vias in routes for a specified number of signals N_{group}. We represent the average number of vias with the quantity $\langle v \rangle$, as given by the expression,

$$\langle v \rangle = \frac{\sum_{i=1}^{N_{group}} v(i)}{N_{group}}, \tag{10}$$

where each sum is taken over all of N_{group} signals in the specified group, and where $v(i)$ is the number of vias in the i^{th} signal in the group. In the case that $\langle v \rangle$ is to be calculated for a group of signals that passes through a specified region, then N_{group} is set equal to the number of signals in that region; in the case that $\langle v \rangle$ is to be calculated for all signals, then $N_{group} = N$, and $\langle v \rangle = v_T/N$.

D.1 Application to the POWER4 IFU

Figures 18- 21 show L_T, $TESL$, $TEML$, and v_T for unit-level interconnections in the DD1 IFU (excluding power and ground) as a function of N_c. The values of the total Steiner length $L_{TS} = \sum_{i=1}^{N} L_S(i)$ and total Manhattan length $L_{TM} = \sum_{i=1}^{N} L_M(i)$ are shown in Fig. 18; L_{TS} is greater than L_{TM} since for each signal i, $L_S(i)$ is an estimate of the required interconnect length based on the design constraints (including macro blockages, which represent tracks that are blocked by the macro from unit-level signal routes, and unit blockages, which represent tracks that are blocked within the unit to reserve these tracks for unit-level signal routes). These figures show that L_T, L_{TS}, L_{TM}, and v_T are reduced for signals targeted for custom interconnections as N_c increases. Moreover, values for these quantities are also reduced for routes that are not targeted for custom interconnections. For $N_c/N = 37\%$, the total interconnect length is reduced by approximately 25mm (0.5%), as shown in Fig. 18. As a result, ΔL_{TS} and ΔL_{TM} are reduced by 13% and 15%, respectively. The total via number is reduced by 20%, as shown in Fig. 21.

Custom interconnections are inserted in seven regions in the DD1 IFU, as shown in Fig. 8. Figure 22 shows $\langle NESL \rangle$ for signals targeted for custom interconnections in (a) the three congested regions (Regions 1, 2, 3) and (b) four uncongested regions (Regions 4, 5, 6, 7) as a function of N_c. For the three congested regions, Fig. 22(a) shows that in Region 1, $\langle NESL \rangle$ remains approximately constant at 0.02 as N_c increases because the overall length of custom interconnections is similar to the overall length provided by the automated router in the wiring bay above the instruction cache stack. Figure 22 shows that $\langle NESL \rangle$ is reduced from 0.02 to 0.01 in Region 2 and from 0.055 to 0.02 in Region 3 for $N_c \geq 3100$. For the four uncongested regions (Regions $4 - 7$) shown in Fig. 8, Fig. 22(b) shows $\langle NESL \rangle$ for the routes targeted for custom interconnections remains unchanged at approximately -0.01 in Region 5 and at -0.05 in Region 6. As shown in

the figure, Region 4 exhibits a large increase in $\langle NESL \rangle$ from -0.4 to 0.3 because the custom interconnections consist primarily of identical routes with L on the order of a few microns. This increase occurs because in Region 4, the automated router routes some signals with zero length but with stacked vias and routes other signals with excessively long routes. Standard design practice recommends that identical, reproducible routes with no stacked vias be used.

Figure 23 shows $\langle v \rangle$ for signals targeted for custom interconnections in (a) the three congested regions (Regions 1, 2, 3) and (b) four uncongested regions (Regions 4, 5, 6, 7), respectively, as a function of N_c. These regions are shown in Fig. 8. For the three congested regions, Fig. 23(a) shows that as N_c increases, the range of $\langle v \rangle$ decreases from $\{6$ to $13\}$ vias per signal to $\{3$ to $6\}$ vias per signal. Note that the apparent dip in the curve for $\langle v \rangle$ that corresponds to Region 1 occurs because the router is unable to complete the routes in this region; therefore, $\langle v \rangle$ is greatly underestimated for these trials. For the four uncongested regions (Regions $4-7$), Fig. 23(b) shows $\langle v \rangle$ for the routes targeted for custom interconnections in Region 5 is decreased by approximately 1 via per signal (from 5 to 4 vias per signal) and that for the targeted routes in Region 4, $\langle v \rangle$ remains nearly unchanged at approximately 3 vias per signal. In Region 6, in which the custom interconnections consist of straight minimum-width $m4$ segments that connect unit-level vertical $m4$ macro bar pins to macro-level horizontal $m3$ 1-bit-wide bar pins, $\langle v \rangle$ is reduced to 1 via per signal when all custom interconnections are added to the design ($N_c = 3401$). In Region 7, in which the custom interconnections consist of straight minimum-width horizontal $m3$ segments that connect several aligned horizontal $m3$ unit-level macro pins, $\langle v \rangle$ is reduced from 2 vias per signal at $N_c = 0$ (as generated by the automated router) to 0 vias per signal at $N_c = 3401$, as expected for signals in which the signal pins are aligned and on the same metal layer; since these signals can be routed with straight metal segments, these signal routes do not require any vias.

VI. Implications for ULSI design

In the preceeding sections, we have seen that interconnect complexity is reduced with each iteration of custom interconnection intervention. However, the net impact of the insertion of custom interconnections on chip performance and yield is still an open issue,

and the influence of decreased interconnect complexity is also unknown.

A. Performance and power

As interconnect length L decreases, interconnect delay τ also decreases, since $\tau = \frac{CR_{\square}L}{2W} + \frac{R_{\square}LC_l}{W}$ where C_l is the capacitive load at the end of the route[8], W is the route width, C is the interconnect capacitance, R_{Box} is the sheet resistance, and the interconnect is modeled as a distributed rc line[36]. Signals with routes on upper-level metal have reduced values of τ, since the values of $R_{\square}$ are lower for these metal layers. The 90% interconnect delay $\tau_{90\%}$ is given by the expression[30],

$$\tau_{90\%} = R \cdot C + 2.3(R_{tr} \cdot C + R_{tr} \cdot C_L + R \cdot C_L), \tag{11}$$

where R_{tr} is the on-resistance of the driver transistor, C_L is the load capacitance on the wire, and R and C scale linearly with L.

The total power P dissipated in interconnections is proportional[31] to the total interconnect capacitance[4], [32], [33], [4], [34],

$$P = \frac{1}{2} \cdot C \cdot V_{dd}^2 \cdot E(sw) \cdot f_{clk}, \tag{12}$$

where f_{clk} is the clock period, and *E(sw)* is the *switching activity* or number of voltage transitions per $1/f_{clk}$ time.

A.1 Application to the POWER4 IFU

Figures 24(a)-(d) show distributions of (a) interconnect netlength L, (b) resistance R, (c) capacitance C, and (d) delay τ, respectively, for all IFU interconnections (excluding power, ground, and clock signal). A two-dimensional extraction of the rc-network of all the unit-level interconnections in a completed physical design (including all metal layers and silicon layers) determines R and C. Each interconnect delay τ is calculated in a simulation of the worst-case static timing of all signal paths under conditions in which the operating voltage $V_{dd} = 1.35$V and temperature $T = 85°$C. Figure 24(a) shows that the DD2 IFU contains fewer signals with long netlengths L. We note that there are a greater number of interconnections with smaller values of R, C, and τ in the DD2 IFU compared with DD1 IFU, as shown in Figs. 24(b)-(d). The peaks in Fig. 24(a) result

from the number of signals that are created from insertion of buffers in unit-level signals with initial unbuffered netlengths greater than $1400\mu m$. The total capacitance of all the interconnections (*ac* interconnections as well as *dc* interconnections) is reduced 3% from 1.173nF in DD1 to 1.136nF in DD2.

The total capacitance of all the interconnections that operate at the full clock frequency (700MHz in DD1; 1.1GHz in DD2) is reduced 3% from 1.08nF in DD1 to 1.05nF in DD2. It follows that the total power P dissipated by switching unit-level interconnections is decreased by 3% from DD1 to DD2.

The largest value of C of all unit-level interconnections is reduced 9.4% from 780.3fF in DD1 to 707.3fF in DD2. These two values occur on two adjacent timing-critical signals on an address bus. In the DD1 IFU, nine signals have values of $C \geq 700$fF, whereas in the DD2 IFU, only one signal wire has $C \geq 700$fF.

Figure 24(d) shows the interconnect delay distribution for τ for all signals including fanout. The number of signals with small values of the interconnect delay increases from DD1 to DD2. Figure 24(d) shows that the worst-case wire delay of all the unit-level signals decreased 17.3% from 204ps in DD1 to 174ps in DD2. These delays occur for two different test signals in DD1 and DD2. The worst-case wire delay decreased 14.7% from 95ps in DD1 to 81ps in DD2 for wires that switch at the full clock frequency. These results show that although the values of the worst-case wire delays are smaller for the DD2 IFU than the DD1 IFU, the portion of the cycle time that is allocated for wire delay is increased more than 30%. This increase occurs for the test signals because the 204ps wire delay occupies 14.3% of the 1430ps DD1 cycle time, and the 174ps wire delay occupies 18.9% of the DD2 920ps cycle time. For the signals in paths that are required to operate at the full clock frequency of 700 MHz in DD1 and 1.1GHz in DD2, the increase occurs because the 95ps wire delay occupies 6.6% of the 1430ps DD1 cycle time, and the 81ps wire delay occupies 8.8% of the DD2 920ps cycle time.

Figures 25 and 26 show visualizations of locations of *ac* signals in the DD2 IFU with the worse timing slews and timing slacks, respectively. In each figure, (a) shows worse falling slew/slack, (b) shows worse rising slew/slack, and (c) shows worse overall slew/slack. Worse overall slews of the *ac* signals are within project requirements (that is, slews are

required to be less than or equal to 330ps compared with the 1ns cycle time.)

A comparison of the relative magnitude of the power P dissipated in unit-level wires with the power dissipated in the macros P_{macros} in DD1 IFU and DD2 IFU is shown in Table V, where for comparison purposes we take f_{clk} = 1GHz, $E(sw) = 0.2$, and V_{dd} = 1.35V. For the calculations of P, the ac load capacitances are 1.08nF and 1.05nF in DD1 and DD2, respectively. A schematic-based power simulation calculates P_{macros}; these simulations assume a 20% switching factor and default capacitance load for all macros and does not include effects of internal parasitics and internal wires[35].

B. Critical area and yield

Chip productivity, or the number of good (functional) chips per wafer, increases as the number of faults per chip decreases, according to the Poisson yield model.[29] Standard design practice recommends interconnections with uniform pattern density.[29], [37]

B.1 Application to the POWER4 IFU

As discussed in a previously, straight regularly-ordered custom interconnections for wide buses generates a uniform signal density n_s across large regions of the IFU, as shown in two $700\mu m$-wide regions in the left-hand side of Fig. 14(a) with 35.6% custom interconnections (green triangles) compared with signal density obtained with 0% custom interconnections (black squares) or 20% custom interconnections (red circles). As a result, the local pattern density is more uniform throughout these regions, and metal thickness variations tend to decrease, which tends to decrease variations in interconnect delay.[37]

C. Noise

To reduce the amount of cross-talk noise between signals in a design, the amount of mutual inductance needs to be limited.[38] This reduction can be achieved with better interconnect design, including interconnections with decreased netlength, increased wire-to-wire spacing, and increased line width, since the amplitude of cross-talk noise increases with wire length, decreases with line spacing, and decreases with line width. Also, crosstalk noise is less in lines manufactured with Cu than in the same lines manufactured with Al.[38]

C.1 Application to the POWER4 IFU

Figure 27 shows visualizations of locations of signals in the DD2 IFU with (a) maximum noise voltage and (b) minimum noise slack at their sinks. The minimum noise slack is the difference between the noise tolerated at each sink and the maximum noise voltage as calculated by the IBM 3DX noise tool. The maximum noise voltage and minimum noise slack are shown in units of volts (V). This figure shows that there are no noise violations for the IFU signals; that is, all signals exhibit noise slack greater than zero.

D. Extendability, design reuse, and schedule

Custom interconnections for timing-critical signals to fix these routes permit easy, fast, and numerous route iterations as required by design changes (i.e., logic changes) in order to meet timing requirements for the remaining signals within the project schedule.

D.1 Application to the POWER4 IFU

Figure 28 shows a progression of the POWER4 IFU floorplan from the first floorplan (GP DD1) to the second (GP DD2). The height and width of each floorplan is shown. The basic sizes and relative arrangements of the macros remain the same in each floorplan. As a result, considerable design time is saved because the custom interconnections can be reused. In particular, the set of custom interconnections which took over 6 months to implement in the DD1 IFU are completely reused with no shorts and no design violations in approximately 3 weeks in the follow-on DD2 IFU. An inspection of the routes in a manufactured chip would show that the routes for some signals (that is, the signals that are routed with custom interconnections) have the same topology, metal wire layer, and via choice in both designs, whereas the routes for the remaining signals (that are routed with the automated router) are different in each design.

VII. Implications of Scaling

One concern for the design of high-performance circuits is that wires are getting relatively slower than devices as the speeds of high-performance transistors continue to increase.[3] To compare the wire delay and circuit delay in different technologies for a fully functional design is to measure the wire delay and circuit delay under similar conditions.

An efficient way to measure the wire delay for a specific route in a design in two technologies is to find a signal, such as a custom interconnection, that has the same route solution in each design. In this context, the phrase *same route solution* indicates that the chosen route should have the same length in unscaled design dimensions prior to implementation in hardware as well as the same metal layers and the same environment (that is, similar adjacent routes). In general, since the majority of the routes in a design are implemented with an automated wiring program, adjacent wires will probably be different for an arbitrary route. However, for the case in which similar routes exist, however, a two-dimensional extraction and a simulation of the worst-case static timing can be performed to obtained the wire resistance R, capacitance C, and delay τ in each technology.

To determine whether wire delays are getting relatively slower or faster compared with devices in each technology, the following test can be performed: Measure the wire delay and device delay in each technology at the same worst-case conditions under which the unit timing is simulated (e.g., 1.35V and 85°C). Devices include NAND-gates, NOR-gates, and inverters[25]. To our knowledge, data for devices measured under worst-case conditions for the DD1 and DD2 versions of the IFU do not exist. To determine whether the relative delay of a wire occupies more or less of the cycle time in different technologies compared to the device delay, we assign the following parameters for the two technologies, as shown in Table VI. First, the cycle times of the two different technologies are denoted by τ_1 and τ_2, respectively. For a selected route r, the wire delay of the route in the two technologies is denoted by τ_{r1} and τ_{r2}, respectively. In this case, we assume that the route length, metal layers, and neighboring routes are equal or as similar as possible and that the delays are measured with equivalent voltage and temperature. Next, for each design, the device delay is denoted by $\tau_{device1}$ and $\tau_{device2}$. For example, for the case in which the device is an inverter, the inverter delay is represented as τ_{inv1} and τ_{inv2}, respectively. To ascertain whether the relative wire delay is greater or less than the device delay in the two technologies, we calculate the amount by which the delays have changed from one technology to the next, relative to the cycle time in each technology. To measure the change in relative delay, we calculate following two parameters: (1) the fractional change in route delay Δ_r and (2) the fractional change in device delay Δ_{device} between the two

technologies, where Δ_r is given by the expression,

$$\Delta_r = \frac{\frac{\tau_{r2}}{\tau_2}}{\frac{\tau_{r1}}{\tau_1}} - 1, \tag{13}$$

and where Δ_{device} is given by the expression,

$$\Delta_{device} = \frac{\frac{\tau_{device2}}{\tau_2}}{\frac{\tau_{device1}}{\tau_1}} - 1. \tag{14}$$

.2 Application to the POWER4 IFU

The IFU contains numerous examples of similar routes in the two technologies discussed in the preceding sections. The cycle times of the two technologies are 1430ps for the DD1 IFU and 910ps for the DD2 IFU. In both DD1 and DD2, eight 32-bit-wide instruction buses are routed over the instruction cache, and their lengths, route layers, and adjacent routes are affected minimally by the decrease in IFU area from DD1 to DD2. The signals are routed primarily with double-wide *m4* for nearly the complete length of the route; the automated router completes the rest of the route. For many of the bus signals, the router completes the route with the same total length in DD2 and in DD1. For the rest of the routes in the buses, the routes are very similar in length. An example of one of these routes r is bit 15 of the bus named *ific-cacheout0* which has length $1318.9\mu m$ in both DD1 and DD2, is implemented nearly completely in *m4* in both DD1 and DD2, and is routed adjacent to power and ground routes in the instruction cache. For this route, the delay is simulated in worst-case conditions (1.35V, 85°C) to be 35ps in both DD1 and DD2. From Eqn. 13, the fractional change in the route delay between the two technologies is seen to be $\Delta_r = \frac{\tau_1}{\tau_2} - 1 = 0.57$. From this value, the relative route delay is seen to increase 57% in DD2 compared with DD1.

VIII. Conclusions

A self-consistent formalism for custom interconnect design is presented for analysis of physical design of on-chip interconnections in ULSI designs. This formalism includes a procedure for custom interconnect design and a procedure to quantify the quality of design interconnections. Application of these techniques to analyze interconnections in the IBM POWER4 Instruction Fetch Unit shows that custom interconnection design optimizes

interconnect quality without adverse affects on route runtime and that time required to reuse custom interconnections in follow-on designs is reduced to a few weeks compared to the several months required for initial implementation. Reduction in total interconnect length and upper-level metal route length has the advantage that additional metal is available to route signals outside the unit when contracts are included in the design process [39]. Implications of these techniques for performance and power; critical area and yield; noise; extendability, design reuse, and schedule; and scaling on ULSI designs is also presented.

Acknowledgments

We thank Gerald Plumb and Byron Krauter, at IBM Server Division in Austin, TX, for their comments on the voltage and temperature of simulations of R, C, and τ. We thank Brian Wallace of Cadence Design Systems in Bellevue, WA, for customer support of the vendor tool Silicon Ensemble. We thank Edward Hughes and Philip Honsinger of IBM Microelectronics in Fishkill, NY, for their explanation of the Steiner algorithm presented in the references. We thank IFU team members, Ed Seymour at IBM Server Division in Austin, TX, and Evanthia Papdopoulou, Zhigang Pan, Leon Stok, Mark Dean, Jim Harper, and Jim Ryan at IBM T. J. Watson Research Center, Yorktown Heights, NY, for discussions. We thank Pong-Fei Lu, Mike Rosenfield, Robert Wisniewski at IBM T. J. Watson Research Center, and Brian Konigsburg at IBM Server Division in Austin, TX, for their comments on the text, and thank Tom Bucelot, Izzy Bendrihem and Kelvin Lewis at IBM T. J. Watson Research Center for technical support. Craig Carter, at IBM Server Division in Austin, TX, wrote the clock route program that is referenced in the text. Manjul Bhushan, at IBM Microelectronics in Fishkill, NY, provided the inverter delay values referenced in the text.

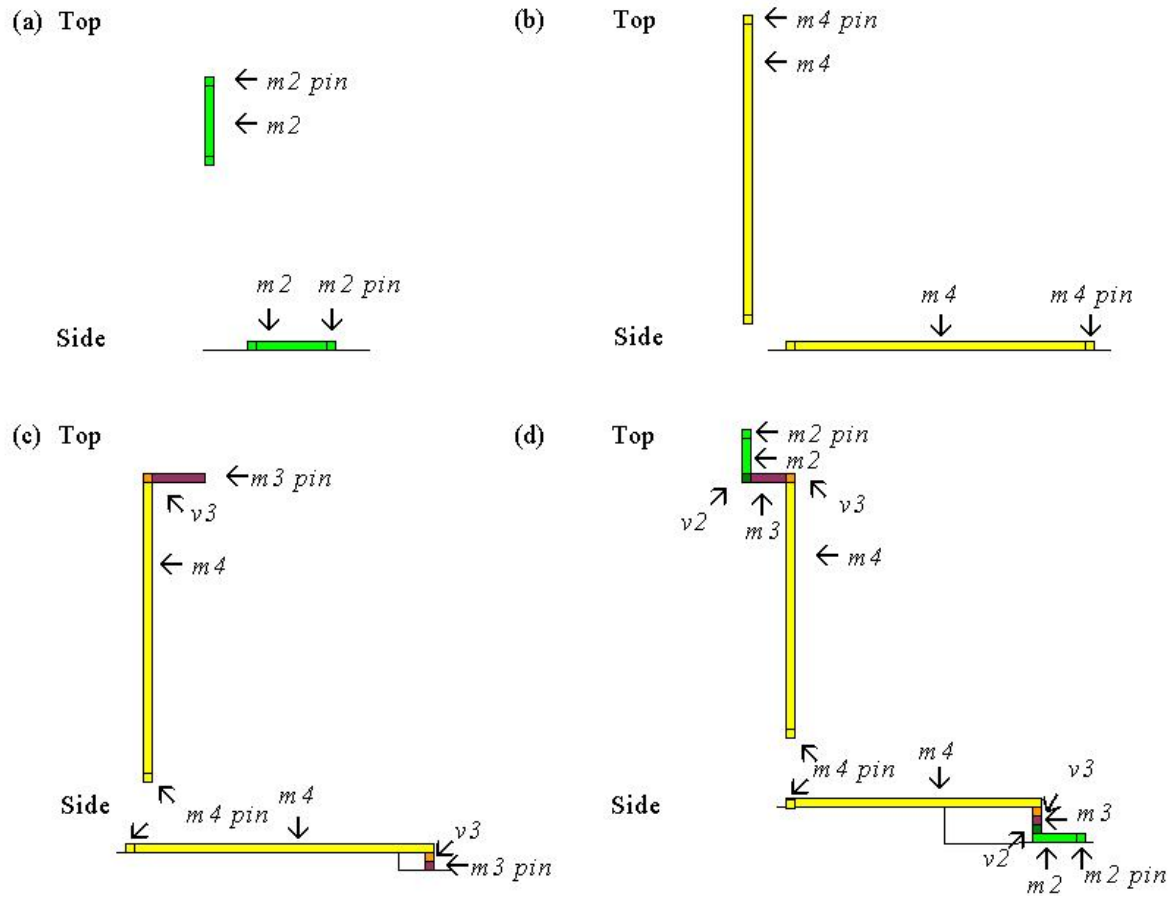

Fig. 1. Examples of custom interconnections for signals with fanout=1. The side and top views for each example are shown. (a) short straight *m2* route between two *m2* square pins that share the same horizontal track on both macros; (b) long straight *m4* route between two *m4* square pins that share the same horizontal track on both macros; (c) straight *m4* route from an *m4* pin to a horizontal *m3* bar pin; (d) straight *m4* route from an *m4* pin to a horizontal *m3* route, then to a vertical *m2* route and *m2* square pin. Vias are also shown in (c) *v3* and (d) *v3*, *v2*.

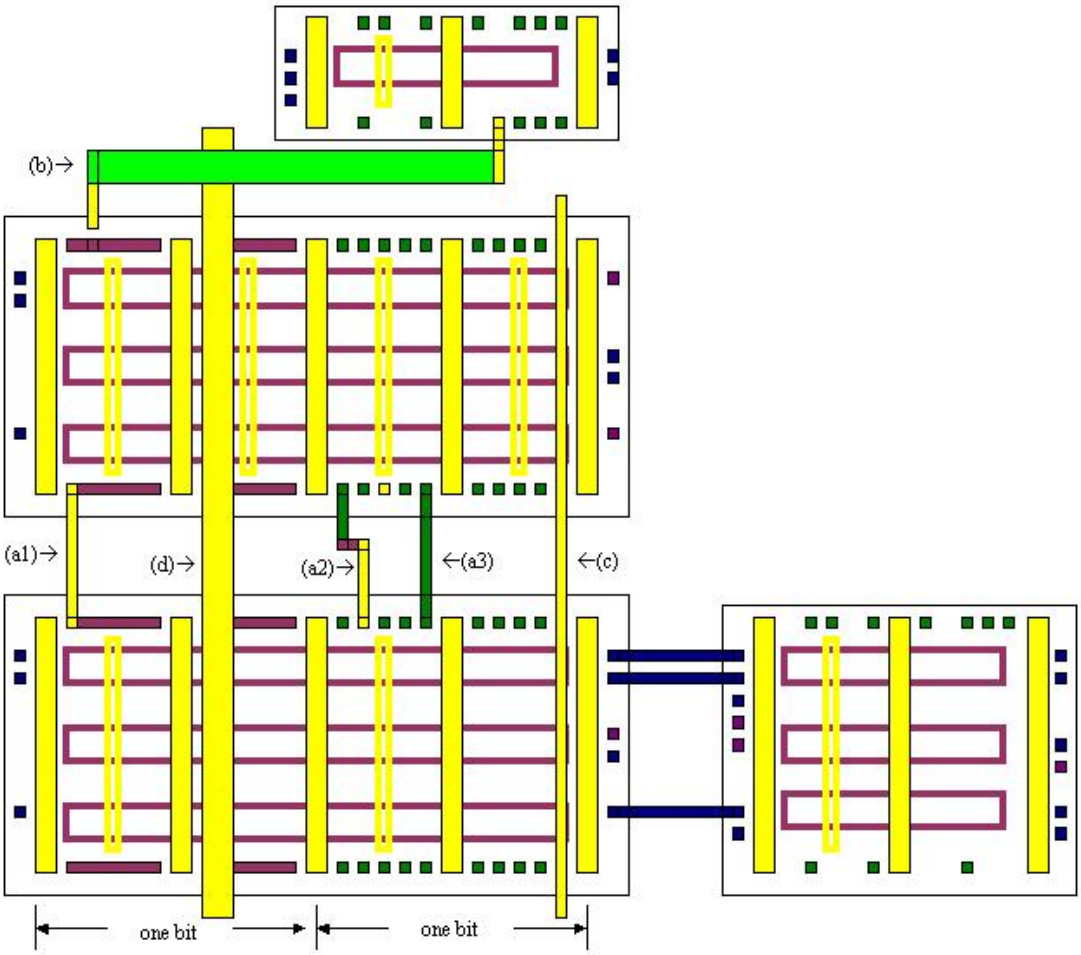

Fig. 2. Examples of different types of custom interconnections in unit-level IFU signal routes. The custom macros and *rlms* are illustrative, since different types of custom interconnections typically occur in different regions of the floorplan instead of on adjacent bits of macros, as shown. The different types of custom interconnections include (*a1-a3*) completely-routed minimum-width routes; (b) completely-routed routes with double-wide routes and minimum-width routes; (c) partially-routed minimum-width routes; and (d) partially-routed wide routes.

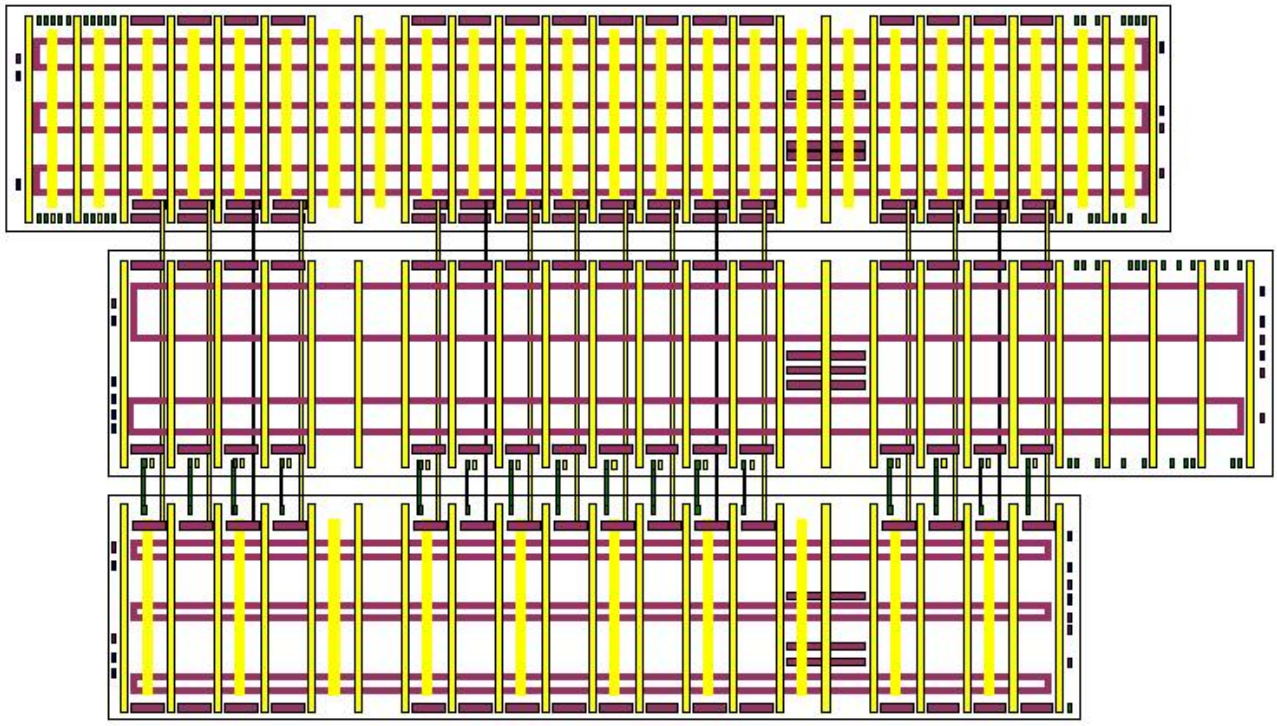

Fig. 3. An example of completely-routed minimum-width custom interconnections to route two 16-bit buses between macros in a dataflow stack. The dataflow stack consists of three 16-bit-wide dataflow macros. In this example, the upper-most macro has a one-bit-wide extension on the right-hand side and a one-bit-wide extension on the left-hand side for control logic. The middle dataflow macro has a two-bit-wide extension on the right-hand-side for control logic, and the bottom dataflow macro does not have extensions. The two two-bit-wide columns and horizontal *m3* pins in these columns represent clock bays and clock control pins, respectively.

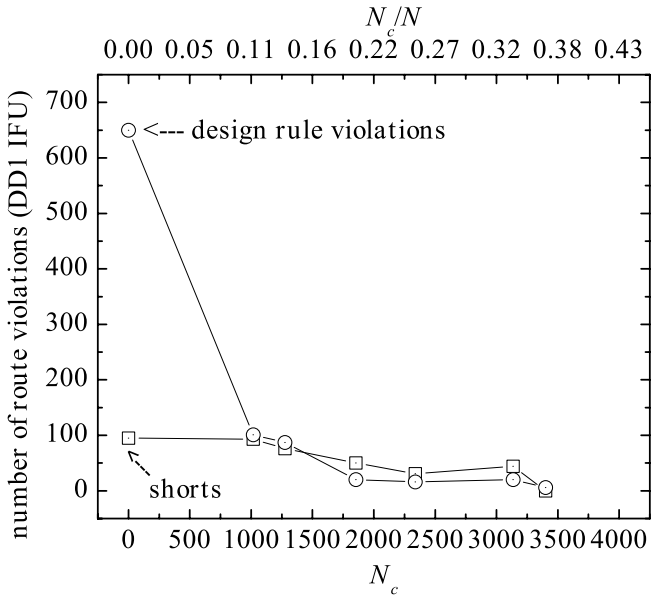

Fig. 4. Number of route violations in the DD1 IFU. Design rule violations (circles) and electrical shorts (squares) are shown as a function of N_c (lower abscissa) and N_c/N (upper abscissa).

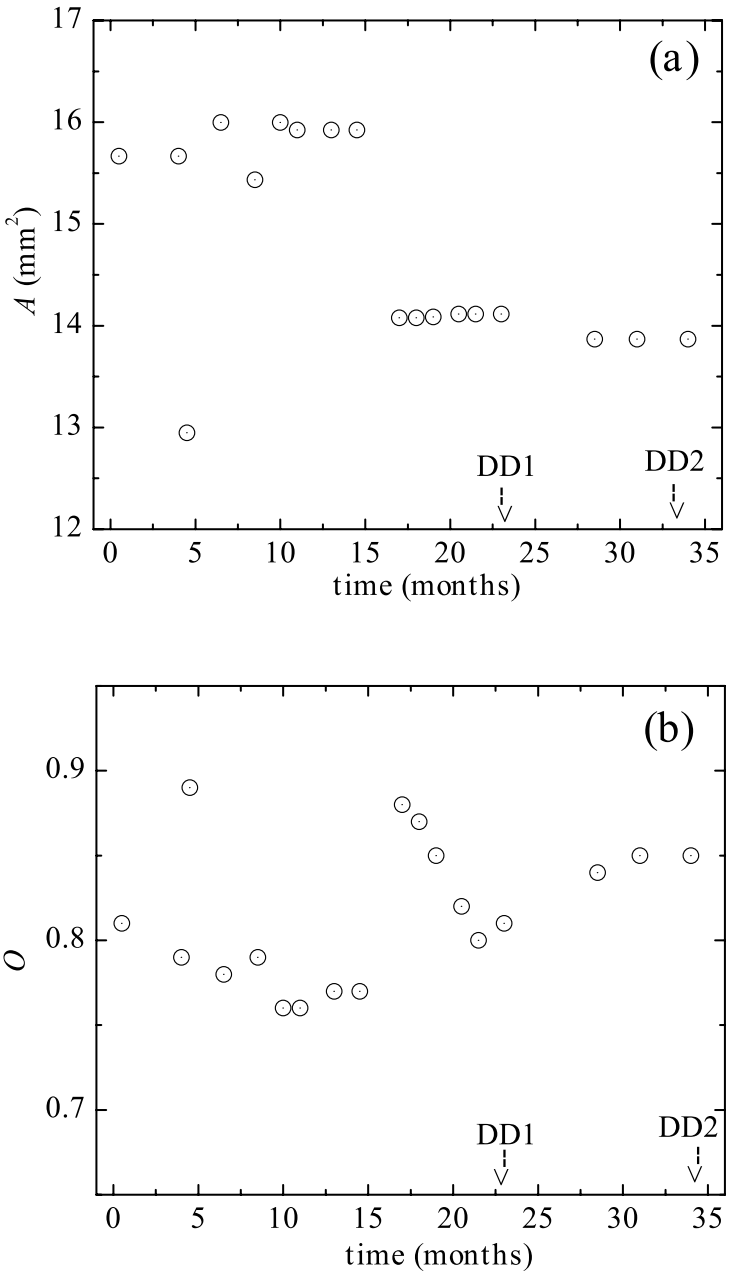

Fig. 5. IFU unit area A (a) and unit occupancy O (b) as a function of time. Snapshots for DD1 and DD2 are indicated at 23 and 34 months, respectively, in both (a) and (b).

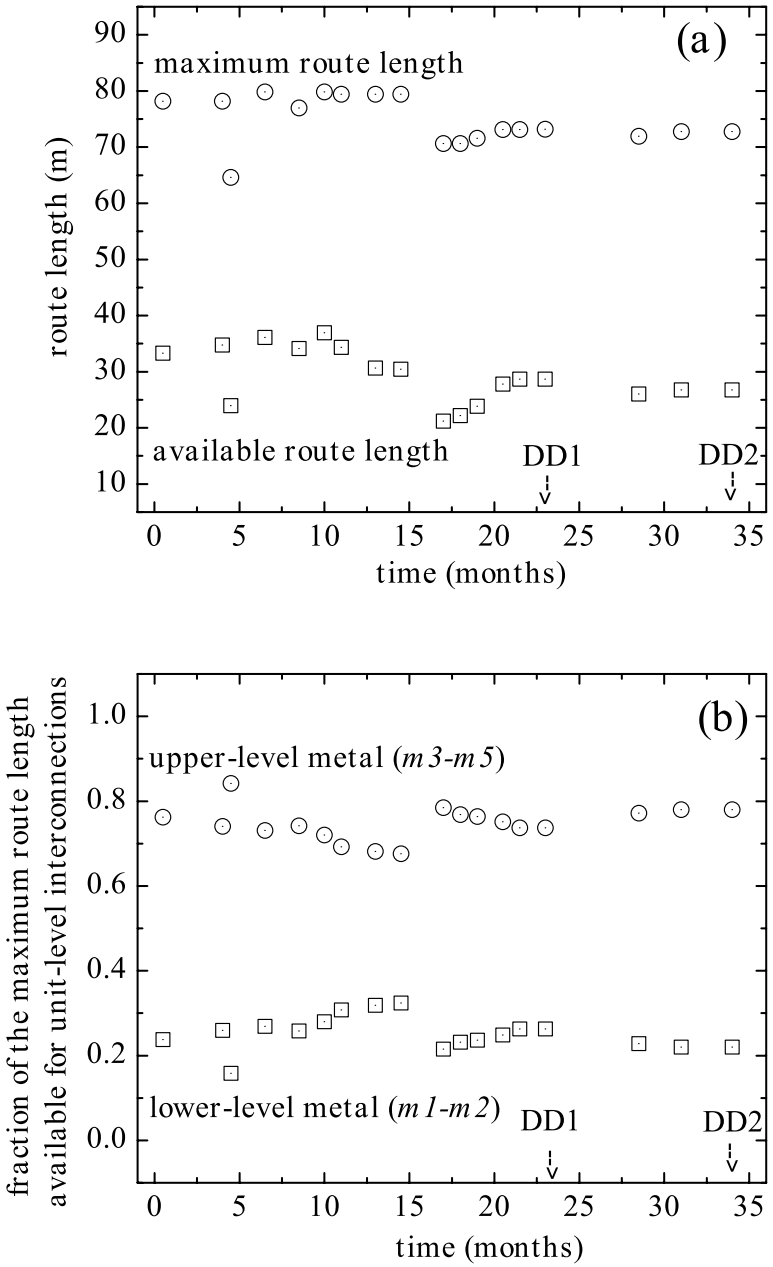

Fig. 6. The maximum route length L_{max} and available route length for minimum-width unit-level signal routes in the IFU as a function of time (a). The fraction of the available route length that is on upper-level metal (*m3-m5*) and on lower-level metal (*m1-m2*) is also shown (b). The values for the DD1 and DD2 snapshots at 23 and 34 months, respectively, are also shown in (a) and (b).

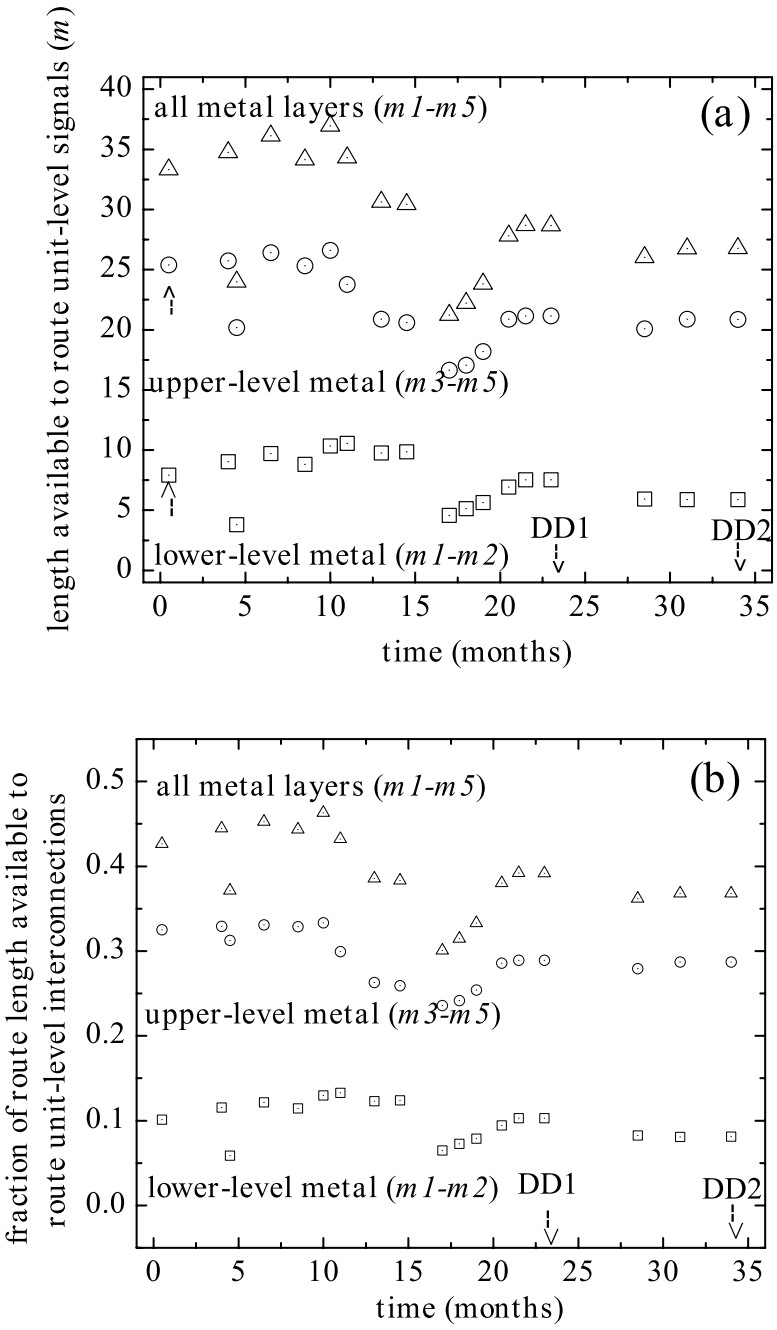

Fig. 7. Interconnect length available to route unit-level signals on each of the five metal layers as a function of time (a). Fraction of the maximum route length that is available to route unit-level interconnections on each of the five metal layers as a function of time (b). The values for the DD1 and DD2 snapshots at 23 months and 34 months, respectively, are shown.

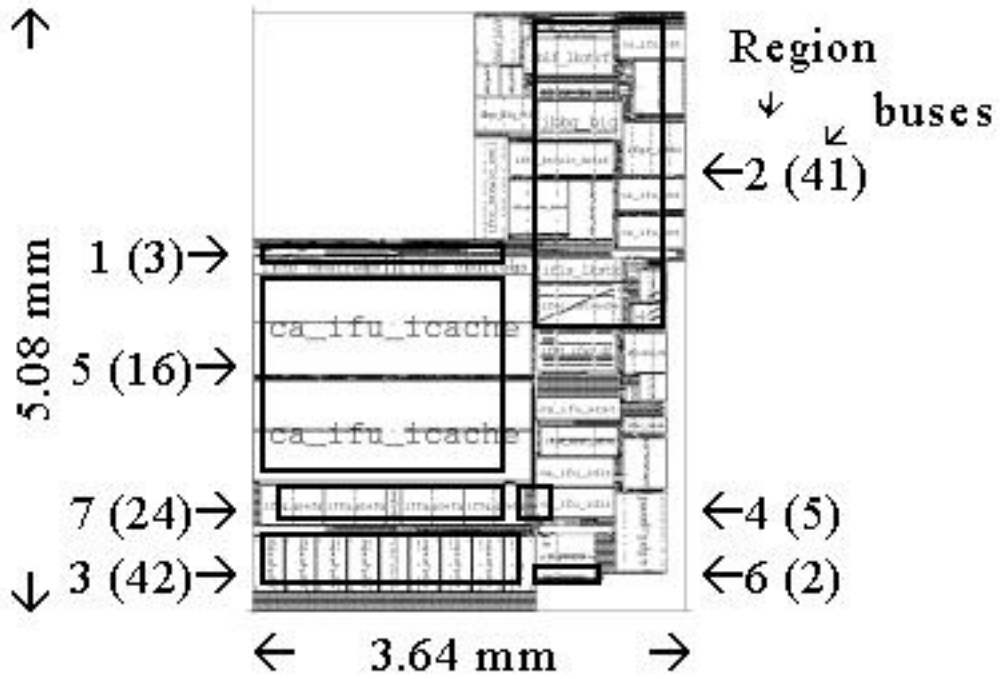

Fig. 8. Regions of the DD1 IFU that are targeted for custom interconnections. The width (3.64mm) and height (5.08mm) of the physical die area are shown. For each region, the number of targeted buses is shown in parenthesis. As discussed in the text, Regions 1, 2, and 3 are congested; Regions $4-7$ are not congested.

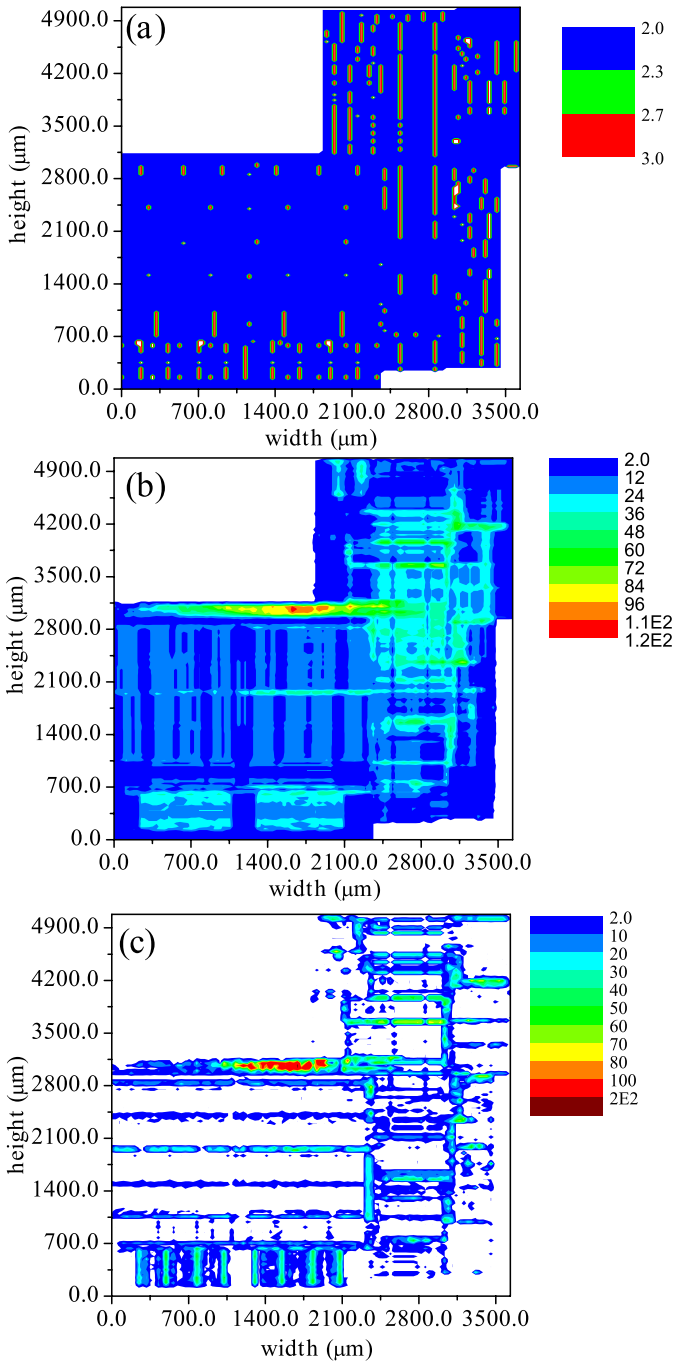

Fig. 9. Visualizations of the signal density n_s and via density n_v in each $35.3\mu m \times 35.3\mu m$ area in the DD1 IFU. The figures show two-dimensional maps of (a) n_s in the unrouted floorplan as well as (b) n_s and (c) n_v in a routed floorplan with no custom interconnections. The power and ground grid sets a lower limit of $n_s = 2$ in the unrouted design (a) except at clock pins, where $n_s = 3$.

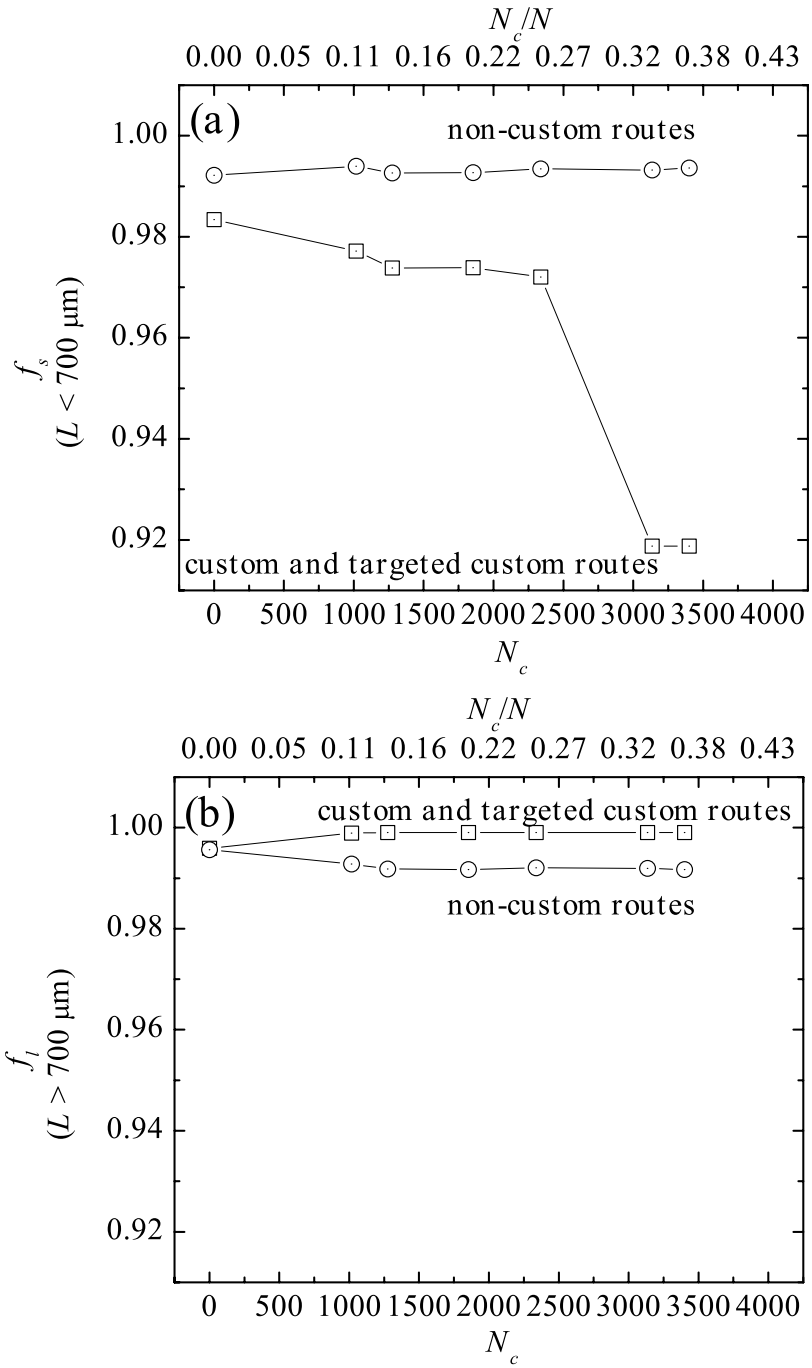

Fig. 10. Fraction of upper-level metal for (a) short (f_s) and (b) long (f_l) signals in DD1 IFU as a function of the number (fraction) of custom interconnections N_c (lower abscissa) and N_c/N (upper abscissa).

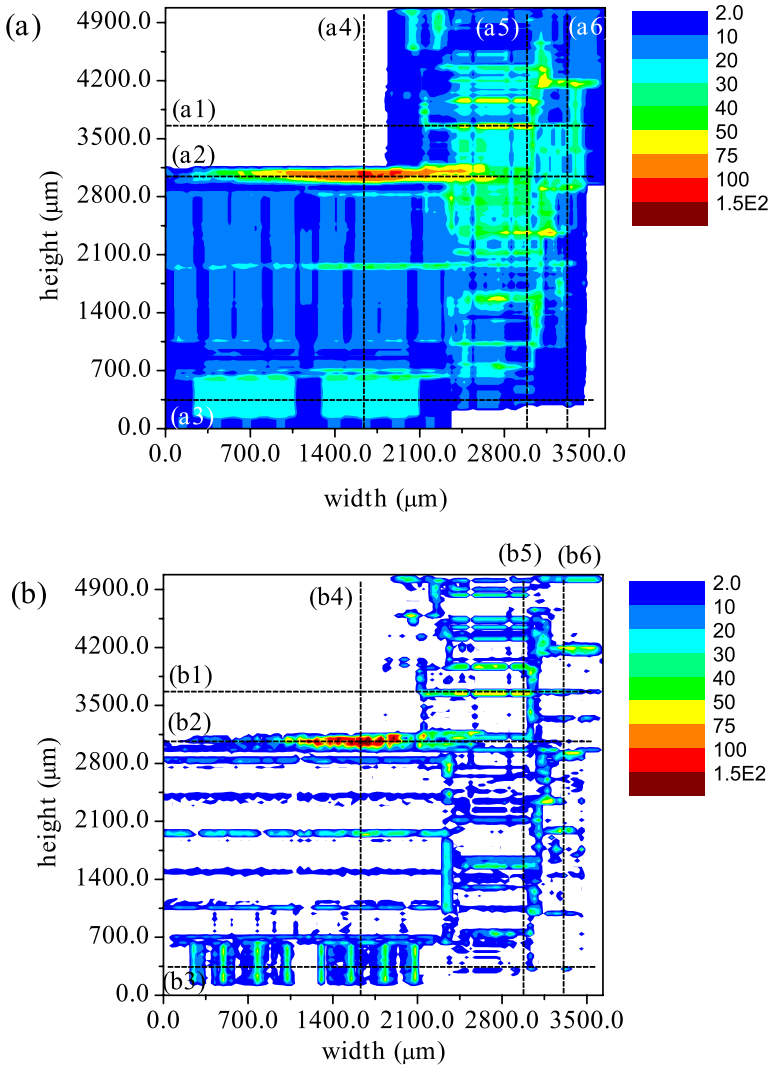

Fig. 11. Visualizations of n_s in a routed design (DD1 IFU) that contains three different values of N_c: (a-a') 0(0%), (b-b') 1853(20.0%), (c-c') 3401(36.8%). (a), (b), and (c) show n_s of custom interconnections; (a') shows n_s in a design routed from (a). (b') and (c') show visualizations of the difference in n_s relative to that in (a') for designs routed from (b) and (c), respectively.

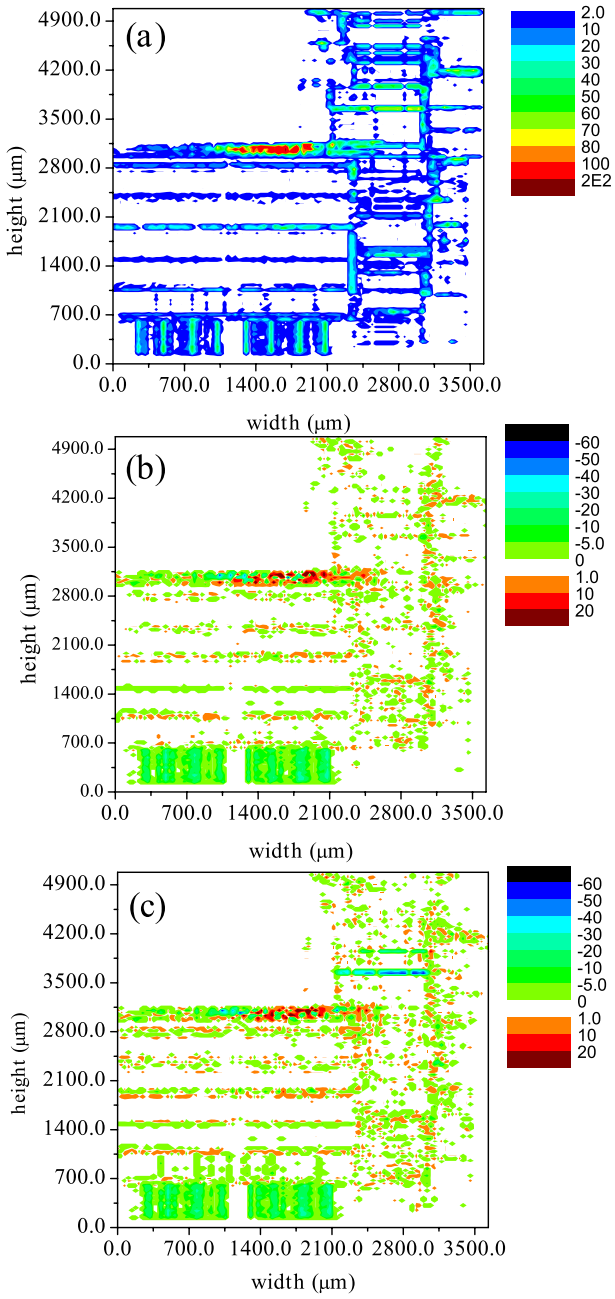

Fig. 12. Visualization of the via density n_v of unit-level interconnections in the DD1 IFU routed with no custom interconnections (a). (b) and (c) show differences in via density relative to that in (a) in which $N_c = 1853$ ($N_c/N = 20.0\%$) and $N_c = 3401$ ($N_c/N = 36.8\%$), respectively.

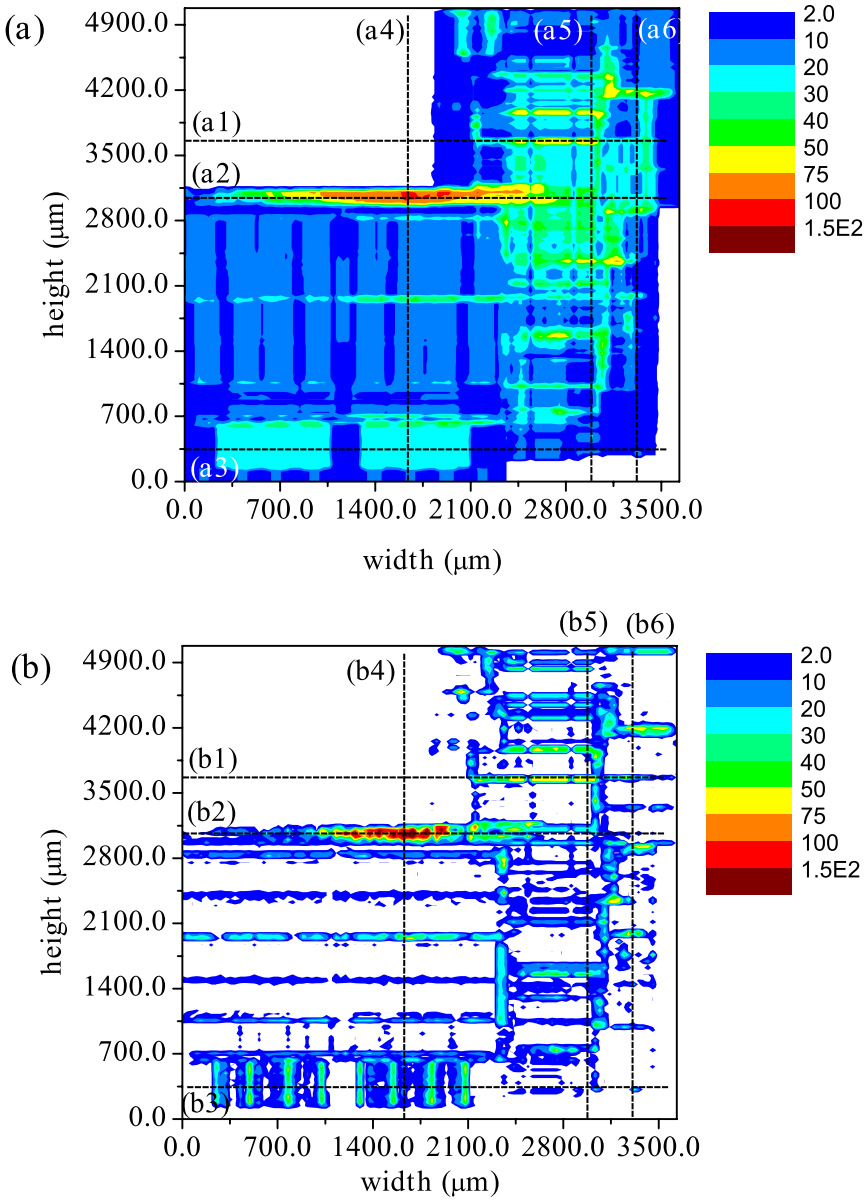

Fig. 13. Visualizations of (a) n_s and (b) n_v per $35.3\mu m \times 35.3\mu m$ area (i.e., 4-bits × 4-bits) in a routed DD1 IFU with no custom interconnections. Dashed lines indicate locations of cross-sections in Fig. 14 and Fig. 15.

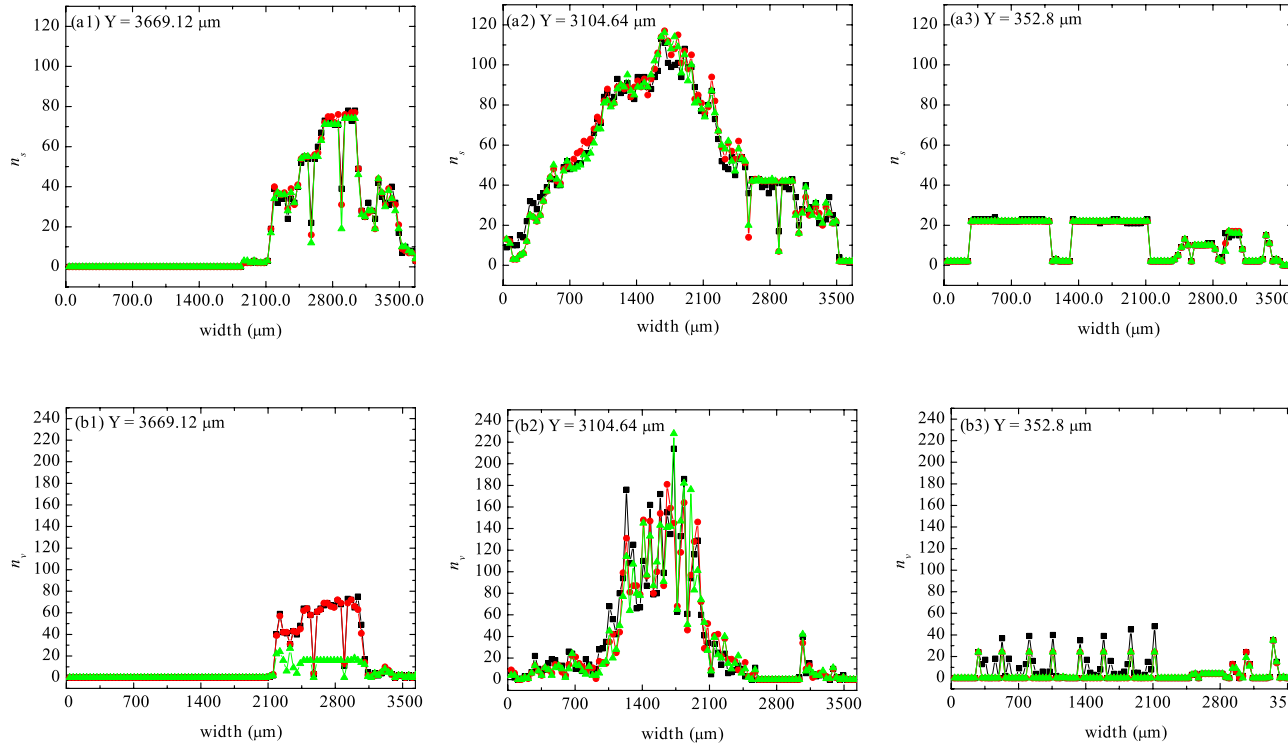

Fig. 14. Horizontal cross-sections of n_s (a1, a2, a3) and n_v (b1, b2, b3) per $35.3\mu m \times 35.3\mu m$ area at three different values of height (1, 2, 3) in a DD1 IFU routed with 0% (black squares), 20% (red cirles), and 35.6% (green triangles) custom interconnections. The cross-sections (a1-a3 and b1-b3) are taken at the locations (dashed lines) shown in Fig. 13.

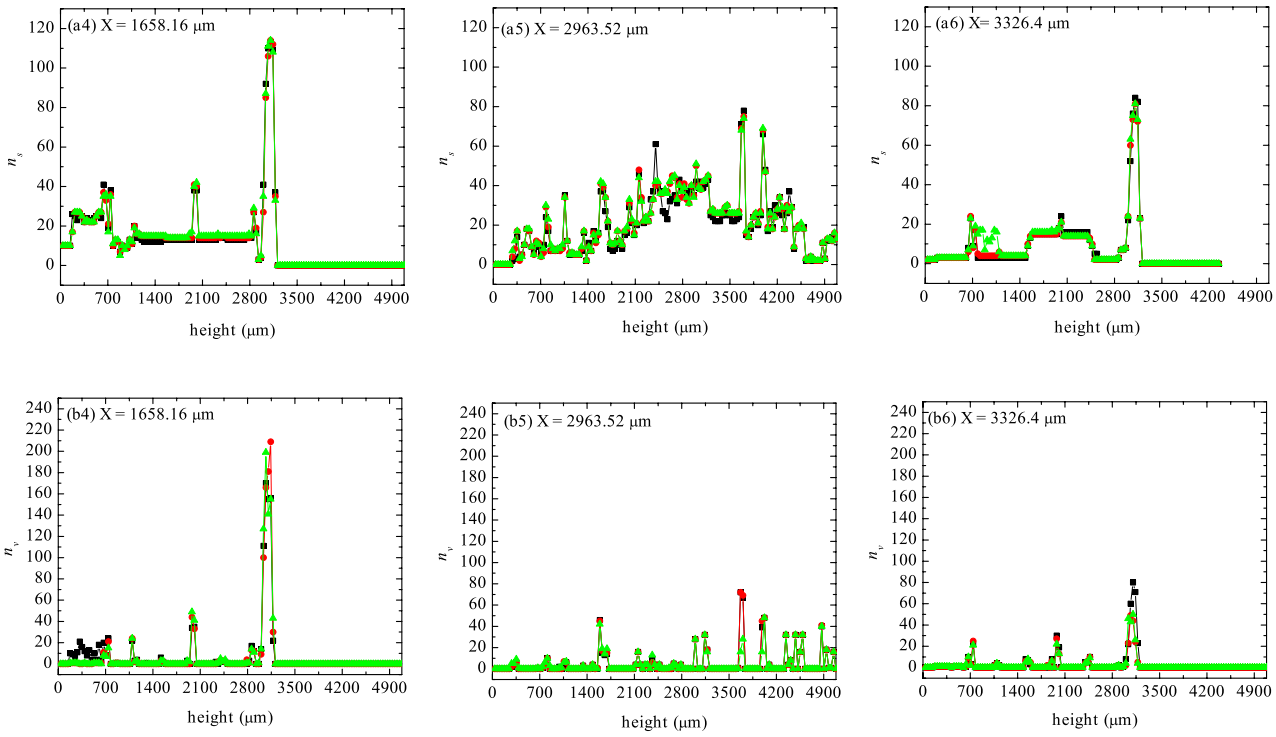

Fig. 15. Vertical cross-sections of n_s (a4, a5, a6) and n_v (b4, b5, b6) per $35.3\mu m \times 35.3\mu m$ area at three different values of the width (4, 5, 6) in a DD1 IFU floorplan routed with 0% (black squares), 20% (red circles), and 38.6% (green triangles) custom interconnections. The cross-sections (a4-a6 and b4-b6) are taken at the locations (dashed lines) shown in Fig. 13.

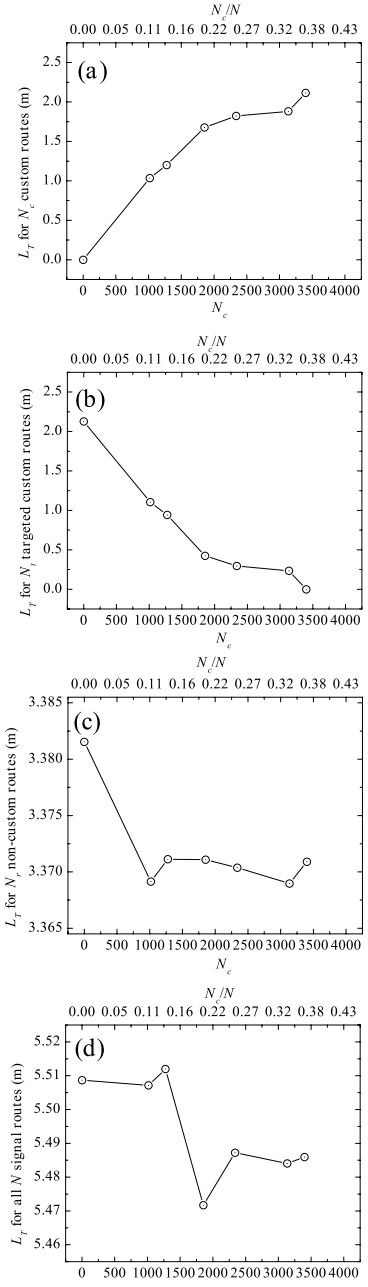

Fig. 16. Total length L_T of (a) custom interconnections, (b) signals that are targeted for custom interconnections, (c) signals with non-custom interconnections, and (d) all N unit-level DD1 IFU interconnections as a function of N_c, lower abscissa (N_c/N upper abscissa), where $N = N_c + N_t + N_r$.

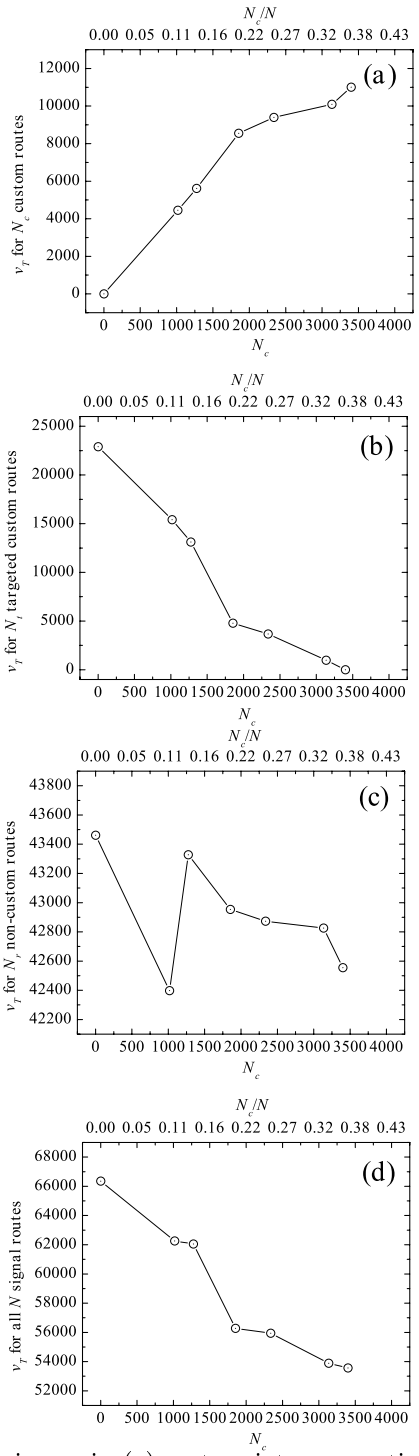

Fig. 17. Total number of vias v_T in (a) custom interconnections, (b) signals that are targeted for custom interconnections, (c) signals with non-custom interconnections, and (d) all N unit-level DD1 IFU interconnections as a function of N_c, lower abscissa (N_c/N, upper abscissa), where $N = N_c + N_t + N_r$.

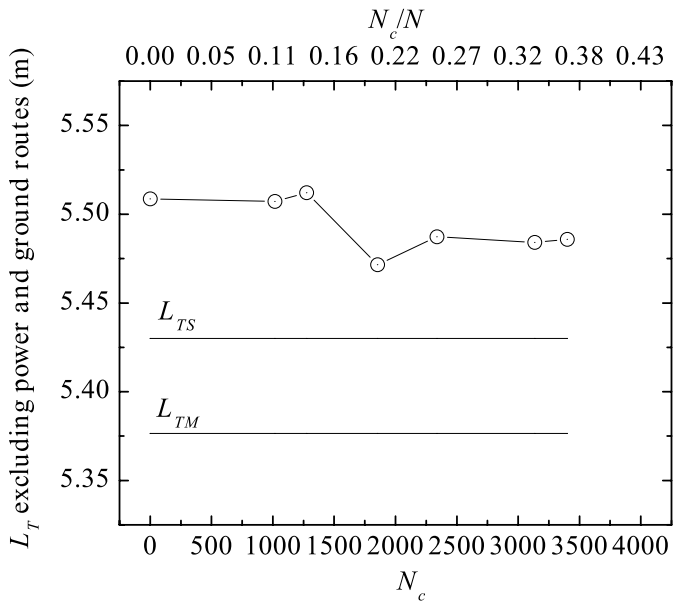

Fig. 18. Total interconnect length L_T of all N DD1 IFU unit-level signals (excluding power and ground) as a function of N_c. The Steiner limit L_{TS} and Manhattan limit L_{TM} are shown.

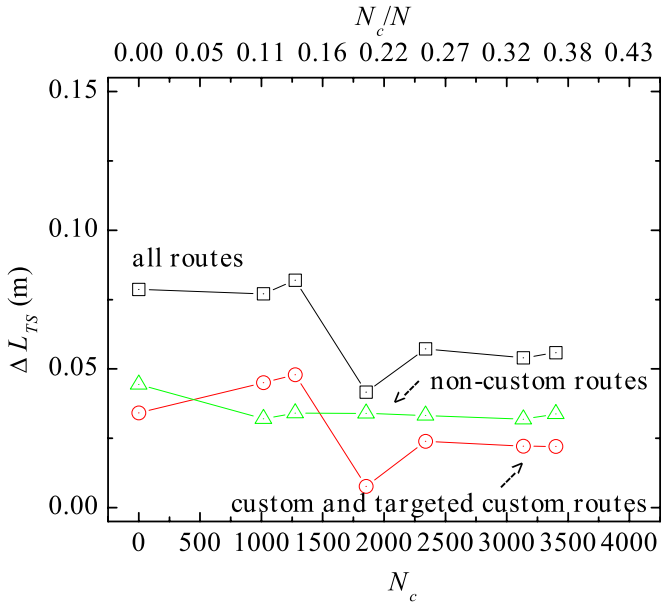

Fig. 19. ΔL_{TS} for all N signals, signals targeted for custom interconnections, and signals with non-custom interconnections as a function of N_c (excluding power and ground routes) in the DD1 IFU.

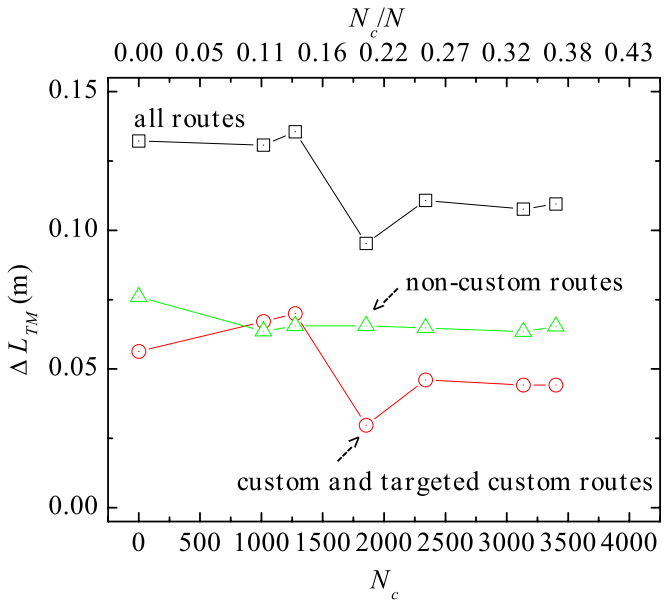

Fig. 20. ΔL_{TM} for all N signals, signals targeted for custom interconnections, and signals with non-custom interconnections as a function of N_c (excluding power and ground routes) in the DD1 IFU.

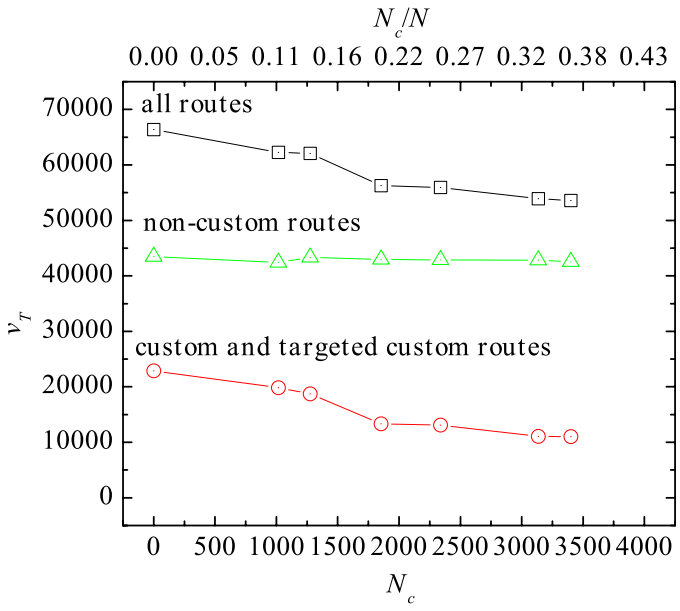

Fig. 21. v_T for all N signals, signals targeted for custom routes, and signals with non-custom interconnections as a function of N_c (excluding power and ground routes) in the DD1 IFU.

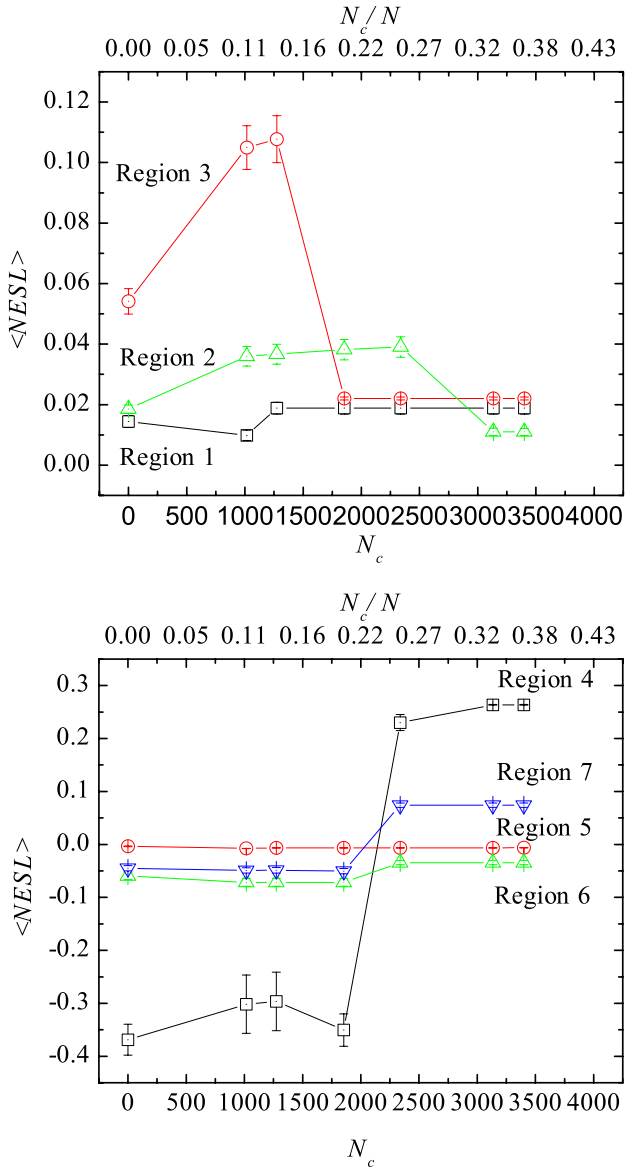

Fig. 22. Average normalized excess Steiner length $\langle NESL \rangle$ for DD1 IFU unit-level signals that are targeted for custom routes as a function of N_c (lower abscissa) and N_c/N (upper abscissa). $\langle NESL \rangle$ is shown for (a) three congested regions (Regions $1-3$) and (b) four uncongested regions (Regions $4-7$). Regions $1-7$ are shown in Fig. 8.

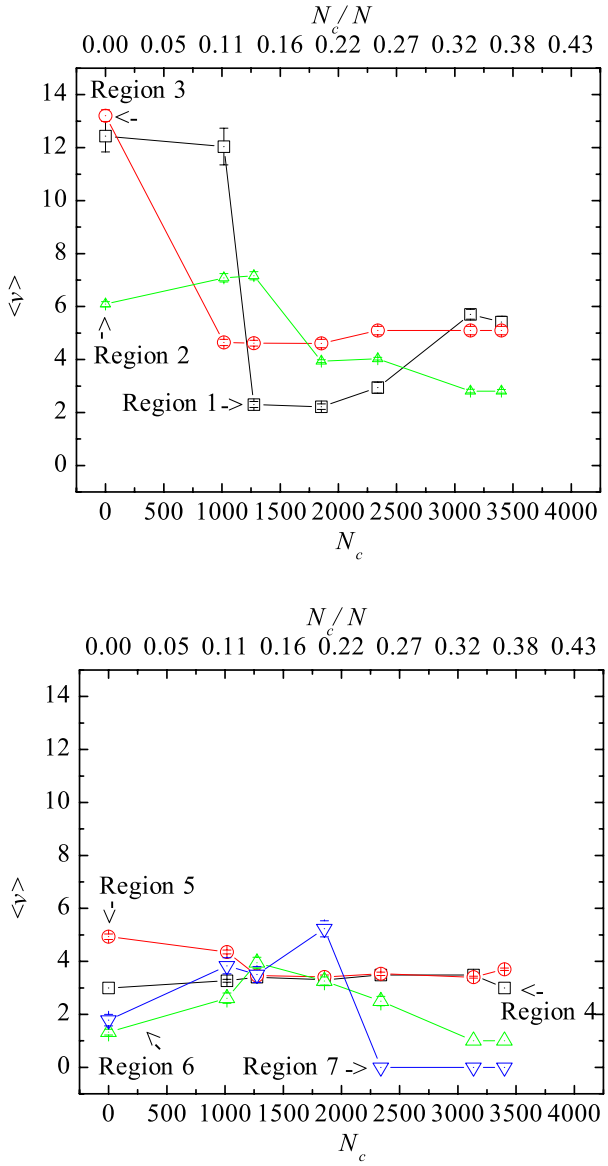

Fig. 23. Average number of vias $\langle v \rangle$ for unit-level signals in the DD1 IFU that are targeted for custom interconnections as a function of N_c (lower abscissa) and N_c/N (upper abscissa). $\langle v \rangle$ is shown for (a) three congested regions (Regions $1-3$) and (b) four uncongested regions (Regions $4-7$). Regions $1-7$ are shown in Fig. 8.

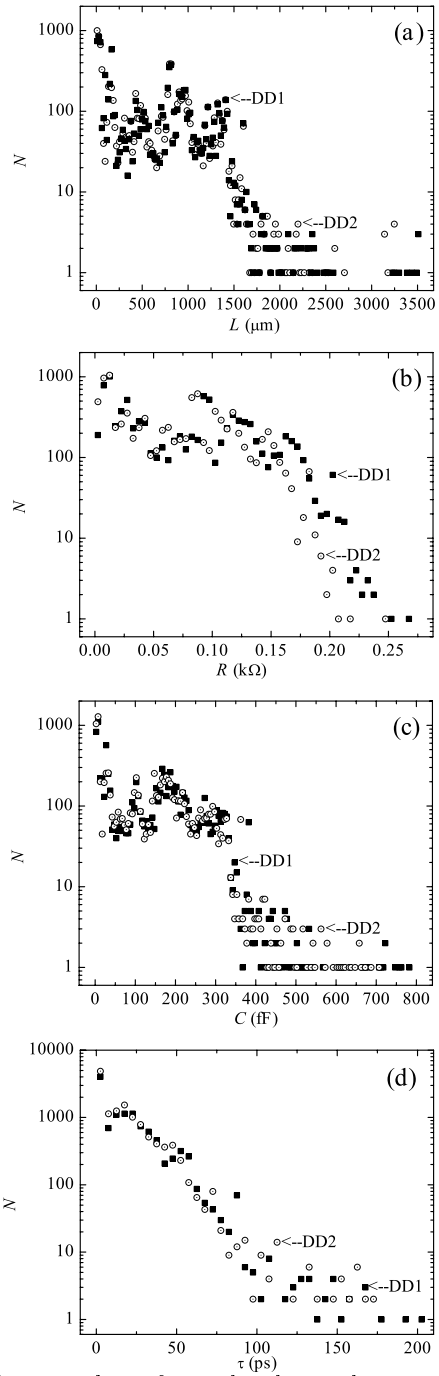

Fig. 24. Distributions of N number of unit-level signals as a function of (a) interconnect length L, (b) interconnect resistance R, (c) interconnect capacitance C, and (d) interconnect delay τ.

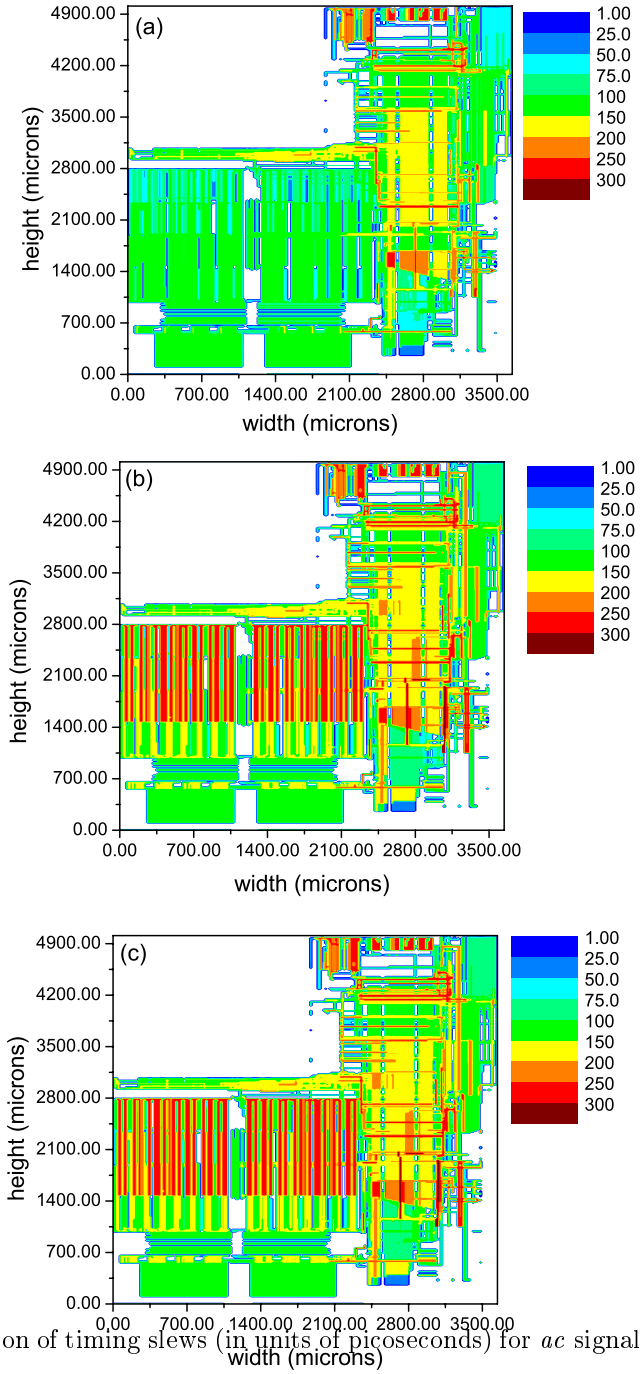

Fig. 25. Visualization of timing slews (in units of picoseconds) for *ac* signal paths in the DD2 IFU. Shown are the locations of *ac* signals with (a) worse falling slew, (b) worse rising slew, and (c) worse overall slew.

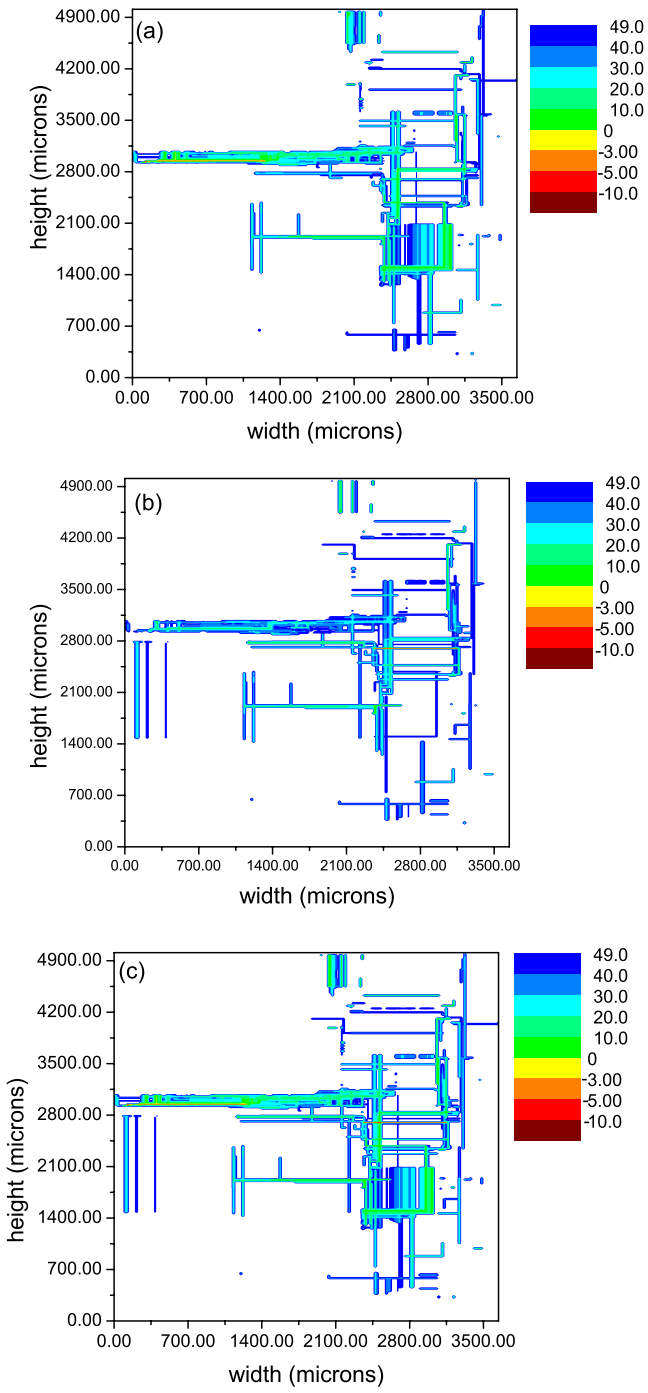

Fig. 26. Visualization of timing slews (in units of picoseconds) for *ac* signal paths in the DD2 IFU. Shown are the locations of *ac* signals with (a) worse falling slack, (b) worse rising slack, and (c) worse overall slack.

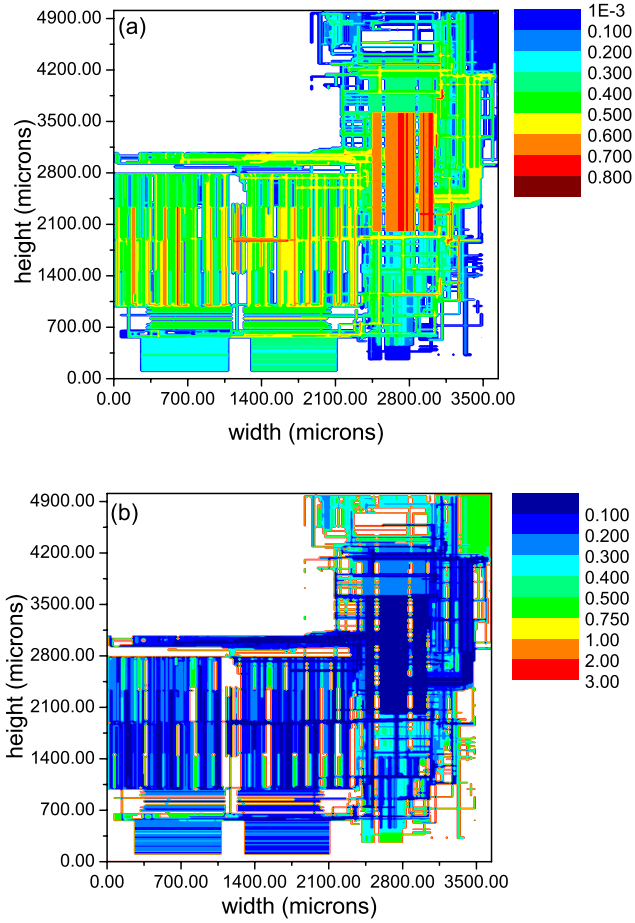

Fig. 27. Visualization of the (a) maximum noise voltage and (b) minimum noise slack on the sinks of signals in the DD2 IFU. The width and height of the IFU are shown in microns, and the noise voltage is in units of Volts. The minimum noise slack is the difference between the maximum noise voltage and the noise that can be tolerated at that sink, as calculated by an IBM noise tool.

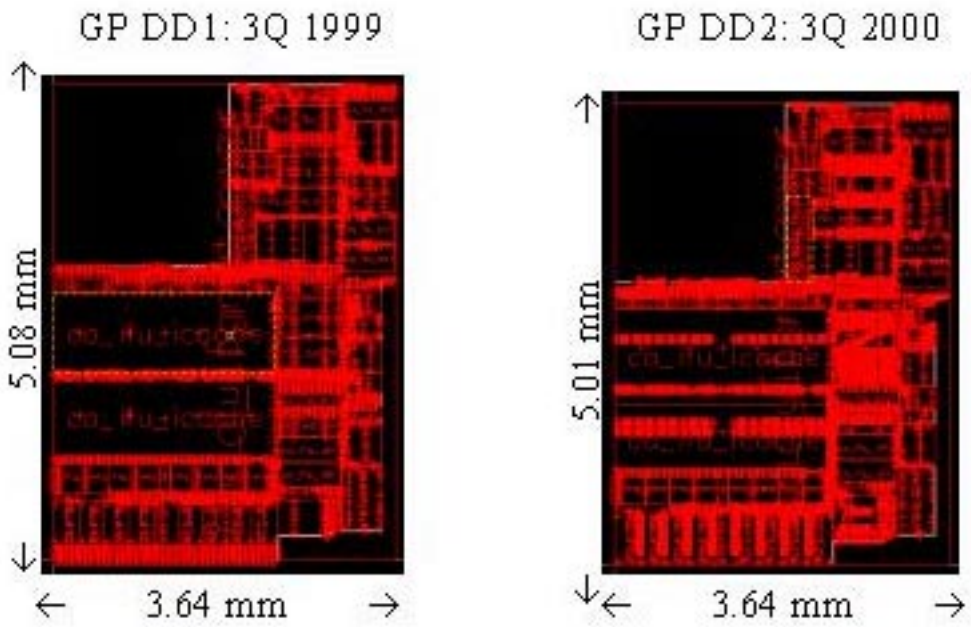

Fig. 28. Progression of the IFU floorplan in GP DD1 (3Q 1999) and GP DD2 (3Q 2000). The width and height of each floorplan is shown.

TABLE I

OVERVIEW OF CUSTOM INTERCONNECTIONS INSERTED IN THE DD1 IFU. FOR EACH TRIAL i, THE NUMBER OF CUSTOM INTERCONNECTIONS N_c^i, NUMBER OF ADDITIONAL CUSTOM INTERCONNECTIONS ΔN_c^i, NUMBER OF BUS SIGNALS N_{bus} AND NUMBER OF CONTROL SIGNALS $N_{control}$ ARE SHOWN. THE TABLE SHOWS THE NUMBER OF BUSES ROUTED IN EACH REGION (N_{reg1} THROUGH N_{reg7}) AND THE TOTAL NUMBER OF BUSES N_{regtot}.

i	N_c^i	ΔN_c^i	N_{bus}	N_{cntl}	N_{reg1}	N_{reg2}	N_{reg3}	N_{reg4}	N_{reg5}	N_{reg6}	N_{reg7}	N_{regtot}
0	0	0	0	0	0	0	0	0	0	0	0	0
1	1016	1016	1016	0	0	19	0	0	0	0	0	19
2	1275	259	248	11	3	0	0	0	8	0	0	11
3	1853	578	578	0	0	0	42	0	0	0	0	42
4	2339	486	486	0	0	0	0	4	0	2	24	30
5	3135	796	796	0	0	22	0	1	1	0	0	24
6	3401	266	266	0	0	0	0	0	7	0	0	7
6	3401	3401	3390	11	3	41	42	5	16	2	24	133

TABLE II

IFU PHYSICAL DESIGN CONSTRAINTS. THE UNIT AREA A, NUMBER N OF UNIT-LEVEL SIGNALS, OCCUPANCY O, AND TOTAL AVAILABLE LENGTH L_{avail} OF MINIMUM-WIDTH WIRE FOR UNIT-LEVEL SIGNALS ARE SHOWN. A DETAILED DISTRIBUTION OF THE UNIT-LEVEL SIGNAL FANOUT IS SHOWN IN THE FOLLOWING TABLE TABLE III.

.

Physical Design Constraints	*DD1*	*DD2*
unit width (μm)	3643.61	3643.61
unit height(μm)	5080.64	5010.08
A (mm^2)	14.1	13.8
number of IO pins	1723	1722
N	9253	9283
average signal fan-out	2.3	2.3
macros (including buffers and inverters)	999	1039
macros (not including buffers and inverters)	95	95
transistors $(\times 10^6)$	5.9	5.9
array transistors $(\times 10^6)$	4	4
buffers and inverters	904	944
O	0.81	0.85
L_{avail} *(meters)*	31.3	28.9
available route fraction	34%	31%

TABLE III

FANOUT DISTRIBUTION OF IFU UNIT-LEVEL SIGNALS IN DD1 AND DD2, EXCLUDING POWER AND GROUND SIGNALS.

.

	Number of Signals	
Fanout	*DD1*	*DD2*
0	328	308
1	7558	7564
2	899	908
3	138	139
4	628	641
5	7	8
6	15	15
7	5	5
8	3	3
95	1	1

TABLE IV

OVERVIEW OF THE NUMBER OF SINGLE STACKED VIAS SSV, DOUBLE STACKED VIAS DSV, TRIPLE STACKED VIAS TSV IN THE DD1 IFU FOR EACH TRIAL i, WHERE THE NUMBER OF CUSTOM INTERCONNECTIONS N_c IN EACH TRIAL IS SHOWN. EACH SSV IS COMPOSED OF TWO STACKED VIAS; EACH DSV IS COMPOSED OF THREE STACKED VIAS; EACH TSV IS COMPOSED OF FOUR STACKED VIAS.

.

i	N_c	SSV	DSV	TSV
0	0	6814	878	18
1	1303	6179	854	12
2	1410	6126	903	9
3	1988	5289	904	10
4	2474	5281	746	16
5	3291	4919	718	14
6	3428	4923	657	12

TABLE V

POWER DISSIPATION IN THE IFU. FOR THE VALUES QUOTED, SWITCHING FACTOR $E(sw) = 0.2$, CLOCK FREQUENCY $f_{clk} = 1\text{GHz}$, $V_{dd} = 1.35\text{V}$ ARE ASSUMED FOR COMPARISON PURPOSES IN BOTH DD1 AND DD2. FOR P, THE ac LOAD CAPACITANCES IN DD1 AND DD2 ARE 1.08NF AND 1.05NF, RESPECTIVELY. SCHEMATIC-BASED POWER SIMULATIONS OBTAIN P_{macros}; FOR THESE SIMULATIONS $E(sw) = 0.20$ WITH A DEFAULT CAPACITANCE LOAD FOR ALL MACROS AND DOES NOT INCLUDE EFFECTS OF INTERNAL PARASITICS AND INTERNAL MACRO WIRES.[35]

	P (Watts)	P_{macros} *(Watts)*
DD1	0.197	3.25
DD2	0.191	3.31

TABLE VI

CHARACTERISTICS OF WIRE AND DEVICE DELAYS IN TWO TECHNOLOGIES. THE CYCLE TIME τ, ROUTE DELAY τ_r AND DEVICE DELAY τ_{device} ARE SHOWN FOR EACH OF TWO TECHNOLOGIES.

Characteristic	*First technology*	*Second technology*
Cycle time	τ_1	τ_2
Route delay	τ_{r1}	τ_{r2}
Device delay	$\tau_{device1}$	$\tau_{device2}$

List of Figures

References

[1] M. T. Bohr, "Interconnect scaling - The real limiter to high performance ULSI," in *IEDM Tech. Dig.*, 1995, pp. 241-244.

[2] A. Deutsch, G. V. Kopcsay, C. W. Surovic, B. J. Rubin, L. M. Terman, R. P. Dunne, Jr., T. A. Gallo, R. H. Dennard, "Modeling and characteristics of long on-chip interconnections for high-performance microprocessors," *IBM J. Res. Dev.* vol. 39, pp. 547-567, Sept. 1995.

[3] ITRS Semiconductor Industry Association, The International Technology Roadmap for Semiconductors, 1999.

[4] R. W. Keyes, "Fundamental limits in digital information processing," *Proc. IEEE*, vol. 69, pp. 267-278, Feb. 1981.

[5] R. W. Keyes, "The wire-limited logic chip," *IEEE J. Solid-State Circuits*, vol. SC-17, pp. 1232-1233, Dec. 1982.

[6] R. W. Keyes, "The power of connections," *IEEE Circ. Dev. Mag.*, vol. 7, pp. 32-35, May 1991.

[7] J. D. Meindl, "Opportunities for gigascale integration," *Solid State Tech.*, pp. 85-89, Dec. 1987.

[8] G. A. Sai-Halasz, "Performance trends in high-end processors," *Proc. IEEE*, vol. 83, pp. 20-36, Jan. 1995.

[9] R. F. Service, "Can chip devices keep shrinking?" *Science*, vol. 274, pp. 1834-1836, Dec. 1996.

[10] P. Solomon, "A comparison of semiconductor devices for high-speed logic," *Proc. IEEE*, vol. 70, pp. 489-509, May 1982.

[11] P. M. Solomon, "The need for low resistance interconnects in future high-speed and high-frequency devices and systems," *Proc. SPIE*, vol. 947, pp. 104-123, Mar. 1988.

[12] A. K. Stamper, "Interconnection scaling 1 GHz and beyond," *IBM MicroNews*, vol. 4, pp. 1-12, 1998.

[13] H. B. Bakoglu and J. D. Meindl, "Optimal interconnection circuits for VLSI," in *ISSCC Dig. Tech. Papers*, 1984, pp. 164-165.

[14] H. B. Bakoglu and J. D. Meindl, "Optimal interconnection circuits for VLSI," *IEEE Trans. Electron Devices*, vol. ED-32, pp. 903-909, May 1985.

[15] W. E. Donath, "Placement and average interconnection lengths of computer logic," *IEEE Trans. Circuits Syst.*, vol. CAS-26, pp. 272-277, Apr. 1979.

[16] D. C. Edelstein, G. A. Sai-Halasz, Y.-J. Mii, "VLSI on-chip interconnection performance simulations and measurements," *IBM J. Res. Dev.* vol. 39, pp. 383-401, Jul. 1995.

[17] M. Feuer, "Connectivity of random logic," *IEEE Trans. Comput.*, vol. C-31, pp. 29-33, Jan. 1982.

[18] Y.-J. Mii, "Performance consideration for the scaling of submicron on-chip interconnections," *Proc. SPIE*, vol. 1805, pp. 332-336, Sept. 1992.

[19] J. Qian, S. Pullela, and L. Pillage, "Modeling the 'effective capacitance' for the RC interconnect of CMOS gates," *IEEE Trans. Computer-Aided Design*, vol. 13, pp.1526-1535, Dec. 1994.

[20] T. Sakurai, "Approximation of wiring delay in MOSFET LSI," *IEEE J. Solid-State Circuits*, vol. SC-18, pp. 418-426, Aug. 1983.

[21] C. J. Anderson, J. Petrovick, J. M. Keaty, J. Warnock, G. Nussbaum, J. M. Tendler, C. Carter, S. Chu, J. Clabes, J. DiLullo, P. Dudley, P. Harvey, B. Krauter, J. LeBlanc, P.-F. Lu, B. McCredie, G. Plumb, P. J. Restle, S. Runyon, M. Scheuermann, S. Schmidt, J. Wagoner, R. Weiss, S. Weitzel, B. Zoric, "Physical design of a fourth-generation POWER GHz microprocessor," in *ISSCC Dig. Tech. Papers*, 2001, pp. 232-233.

[22] K. Diefendorff, "Power4 focuses on memory bandwidth," *Microprocessor Report*, vol. 13, pp. 1-8, Oct. 1999.

[23] IBM, NY. IBM Enterprise server pSeries 680. [Online] Available: http://www-1.ibm.com/servers/eserver/pseries/hardware/enterprise/

[24] J. D. Warnock, J. Keaty, J. Petrovick, J. Clabes, C. J. Kircher, B. Krauter, P. Restle, B. Zoric, C. J. Anderson, "The circuit and physical design of the POWER4 microprocessor," *IBM J. Res. Dev.*, vol. 46, pp. 27-51, Jan. 2002.

[25] M. Bhushan, private communication, Nov. 28, 2000. One way to characterize IBM's technologies is to measure the delay of an unloaded inverter (typically, with a ring oscillator) operated at 1.5 Volts and 25 ° C. (Note that these conditions are different than those cited in our paper because the hardware is typically measured under more advantageous conditions, i.e., higher voltage and lower temperature, than the conditions under which the design is simulated, i.e. worst-case conditions with lower voltage and at higher temperature). For the DD1 technology, M. Bhushan measured the delay of a single fanout inverter to be 14.5-16 ps, and the delay of an inverter with a fanout of three to be 30-33 ps. For the DD2 technology, M. Bhushan measured the delay of a single fanout inverter to be 12-13.5 ps, and the delay of an inverter with a fanout of three to be 25-28 ps.

[26] Edward Hughes and Philip Honsinger, personal communication. The Steiner route length of a signal in a circuit design is the complete specification of all the lengths of straight horizontal and vertical route segments that are estimated to connect the points (pins) on a signal. The Steiner problem is NP-complete, which means that the time required to determine an exact route solution for a signal that connects an arbitrary number of pins increases exponentially as the number of pins increases. The Steiner algorithm discussed in this paper is available in IBM's Hierarchical Design Planner.

[27] The Manhattan length L_M is the sum of the maximum horizontal separation of all the macro pins on a signal and the maximum vertical separation of all the macro pins on the same signal. For minimum area square pins, the center of the pin is taken in the calculation. However, for cases in which a signal connects minimum-width square pins to very wide pins, and the wide pin has the largest horizontal separation to the rest of the pins, the track of the wide pin that is nearest the rest of the pins is entered in the calculation of the Manhattan distance L_M. The calculation of the Manhattan length does not take blocked tracks into account in the calculation. In addition, the value of the Manhattan length can be much smaller than the minimum physical distance required to route a signal..

[28] Craig Carter wrote the program that routes the clock routes in the DD1 and DD2 IFU floorplans.

[29] *IBM CMOS Design Manual School*, Nov. 1996.

[30] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI.* New York, NY: Addison-Wesley, 1990.

[31] Rule of thumb. This is true for the IFU interconnections since the lengths of IFU interconnections satisfy the rule of thumb for a wire to be modeled by a capacitance.[30] According to this rule of thumb, the time t_{rf} required by a signal to rise or fall from 10% to 90% of its final value should be greater than 2.5 times the signal time-of-flight t_{flight}, as described by the expression,[30]

$$t_{rf} \geq 2.5 t_{flight} = 2.5 \frac{L}{v}, \tag{15}$$

where $v = c_o/\sqrt{\epsilon}$ is the propagation speed of the signal in the interconnect, c_o is the speed of light, and ϵ is the dielectric constant. For the IFU interconnections, the dielectric constant $\epsilon \sim 4.2$, and the signal propagation speed v is then ~ 14.6 cm/sec. The longest route in the DD2 IFU has $L \sim 3300$ microns, for which Eqn. 15 shows that t_{rf} should be greater than 57 ps. Since typical values of rise-times and fall-times for IFU interconnections range from $\sim$ 100 ps to 330 ps, Eqn. 15 is satisfied for the IFU interconnections..

[32] B. Krauter, S. Mehrotra, and V. Chandramouli, "Including inductive effects in interconnect analysis," in *Proc. CICC*, 1999.

[33] M. Pedram, "Tutorial and survey paper: Power minimization in IC design: Principles and applications," *Trans. Design Automation of Electronic Syst.*, vol. 1, pp. 3-56, Jan. 1996.

[34] J. P. Uyemura, *Circuit Design for CMOS VLSI.* Norwell, MA: Kluwer Academic Publishers, 1992.

[35] P.-F. Lu, private communication, 2001.

[36] Jan M. Rabaey, "Digital Integrated Circuits: A Design Perspective. Upper Saddle River, NJ: Prentice Hall, 1996.

[37] T. Park, T. Tugbawa, D. Boning, J. Chung, S. Hymes, R. Muralidhar, B. Wilks, K. Smekalin, and G. Bersuker, "Electrical characterization of copper chemical mechanical polishing," unpublished.

[38] G. Servel, F. Huret, E. Paleczny, J. F. Legier, D. Deschacht, "Inductive effects on crosstalk evaluation," in *Proc. IITC*, 2001, pp. 63-65.

[39] K. Barkley and T. Wohlfahrt, "Hierarchical methodology: Contract generation and layout abstraction," in *Proc. ICUGC*, 1999.

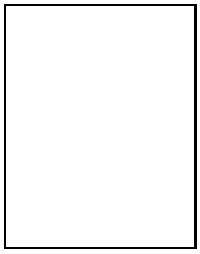

Mary Y. L. Wisniewski received an A.B. summa cum laude from Harvard-Radcliffe Colleges, an M. Phil. from Cambridge University, and a Ph.D. in 1997 from Cornell University, all in Physics. In 1996, she joined IBM Research, Yorktown Heights, NY, as a Research Staff Member in the T. J. Watson Research Center. She is a member of the VLSI Design Department and has contributed to the circuit design of an IBM S/390 microprocessor and was integrator of the Instruction Fetch Unit of the IBM pSeries POWER4 microprocessor. She is currently involved in CMOS VLSI design.

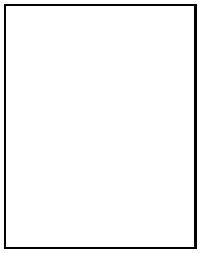

Emmanuel Yashchin received a Diploma in Applied Mathematics from Vilnius State University (U.S.S.R.) in 1974, a M.Sc. in Operations Research and the D.Sc. in Statistics from the Technion-Israel Institute of Technology in 1977 and 1981, respectively. In 1982 he was a Visiting Assistant Professor at the Iowa State University. He joined IBM in 1983 as a Research Staff Member and is manager of the Statistics Group in the Mathematical Sciences Department at the IBM T. J. Watson Research Center, Yorktown Heights, NY. His research interests include Change Point Models, Quality Control, Reliability, Spatial Statistics, Extreme Value Theory and Operations Research.

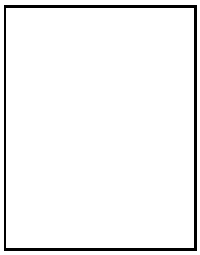

Robert L. Franch received the B.S.E.E. degree from the Polytechnic Institute of New York, Brooklyn, in 1980. He joined IBM East Fishkill, NY in 1980 where he worked on bipolar device reliability and stress testing of VLSI devices and interconnects. In 1984 he joined IBM Research as a member of the Test Systems Group where he worked on high speed functional testing of NMOS and CMOS chips developed at IBM Research. In 1993 he joined the VLSI Design Group where he has worked on circuit design for microprocessors. He has since been engaged in the design of CMOS circuit macros and has contributed to several generations of IBM S/390 microprocessors.

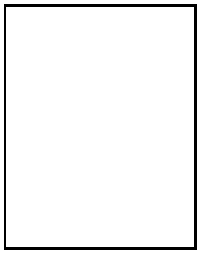

Giovanni Fiorenza received the B.S., M.S., and Ph.D. degrees in Physics from Stevens Institute of Technology in 1979, 1984, and 1992 respectively. He joined the U.S. Air Force 1979 and worked at the Air Force Weapons Laboratory, Kirtland AFB, Albuquerque, NM, on Laser Damage in Thin Films until 1981. In 1984, he joined IBM Component Vendor Assurance, Poughkeepsie, NY, where he worked on technology assessment, functional reliability, and soft-error evaluations of NMOS and CMOS technologies. In 1988, he joined IBM East Fishkill, NY, where he worked until 1992 on the development of high performance BIPOLAR/BICMOS technologies, and until 1996 on the development of device simulator FIELDAY. In 1996, he joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he is currently working on CMOS VLSI design.

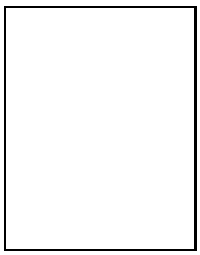

I. Cevdet Noyan works at the IBM T. J. Watson Research Center in Yorktown Heights, NY.